Supervised Learning in Spiking Neural Networks with FORCE Training

Wilten Nicola and Claudia Clopath

November 11, 2016

arXiv:1609.02545v2 [q-bio.NC] 10 Nov 2016

Abstract

Populations of neurons display an extraordinary diversity in the types of problems they solve and behaviors they display. Techniques have recently emerged that allow us to create networks of model neurons that solve tasks of similar complexity. Examples include the FORCE method, a novel technique that harnesses chaos to perform computations. We demonstrate the direct applicability of FORCE training to spiking neurons by training networks to mimic various dynamical systems in addition to reproduce more elaborate tasks such as input classification, storing sequences, reproducing the singing behavior of songbirds, and recalling a scene from a movie. Post-training network analysis reveals behaviors that are consistent with electrophysiological data, such as the stereotypical decrease in voltage variance upon input presentation, reproducing firing rate distributions from songbird data, and reproducing locations of incorrect recall in sequence replay. Finally, we demonstrate that theta oscillations are critical for both learning and recall of episodic memories.

Introduction

Human beings can naturally learn to perform a wide variety of tasks very quickly and efficiently. Examples include learning the complicated sequence of motions in order to take a slap-shot in Hockey, learning the particular routines of our friends and family, or learning to replay the notes of a song after music class. While there are a wide variety of approaches to solving these different problems in various fields such as machine learning, control theory, etc., human beings use a very distinct tool to solve these problems: the spiking neural network.

Our attempts at understanding how these populations of neurons solve problems can be broadly assigned to one of two categories: bottom-up or top-down. Bottom-up approaches start with a biologically constrained network of model neurons. They attempt to determine what the network function is or how a particular property of the network behavior emerges (e.g. [Coombes, 2005, Bressloff, 2011, Knight, 1972, Wilson and Cowan, 1972, Nesse et al., 2008, Omurtag et al., 2000b, Nykamp and Tranchina, 2001, Sirovich et al., 2006, Strogatz and Mirollo, 1991, Strogatz, 2000, Omurtag et al., 2000a, van Vreeswijk, 1996, Van Vreeswijk et al., 1994, Abbott and van Vreeswijk, 1993]). Broadly speaking, given a spiking neural network, a bottom-up approach may tell us what particular problem a network is trying to solve.

Bottom-up approaches can yield substantial insights into a networks behavior, however, they give us no control over this behavior. Top-down approaches, however allow us to enforce a certain behavior or dynamics onto a neuronal network [Sussillo and Abbott, 2009, DePasquale et al., 2016, Abbott et al., 2016, Thalmeier et al., 2015, Boerlin et al., 2013, Schwemmer et al., 2015, Bourdoukan and Denève, 2015, Eliasmith et al., 2012, Eliasmith and Anderson, 2004]. These techniques start with a task or dynamical behavior that a recurrent spiking neural network should perform, and determine what the connection strengths between neurons should be to achieve this. The dominant approaches are currently the FORCE method [Sussillo and Abbott, 2009, DePasquale et al., 2016], spike-based or predictive coding networks [Boerlin et al., 2013, Schwemmer et al., 2015, Bourdoukan and Denève, 2015], and the Neural Engineering Framework (NEF) [Eliasmith et al., 2012, Eliasmith and Anderson, 2004].

While the NEF and spike-based coding techniques can reproduce a wide variety of dynamics in spiking neural networks, they are limited in two very important ways. First, these approaches are not agnostic towards the underlying network of neurons. The weight matrix solutions are derived assuming specific types of synapses or neurons and are not easily extended when one adds an increasing degree of realism to the network or changes the underlying neurons or synaptic connection types. Second, in order to apply either approach, one has to know the underlying dynamics of the task to be performed by the network in terms of closed-form differential equations. This is a restrictive constraint on the potential problems these networks can solve. In

other words, these approaches are limited in their applicability to spiking neuronal networks and limited in their ability to solve problems using these networks.

Fortunately, the FORCE method is agnostic towards both the network under consideration, or the tasks that the network has to solve. This technique falls under the realm of reservoir computing in machine learning [Lukoševičius et al., 2012, LukošEvičIus and Jaeger, 2009, Schrauwen et al., 2007]. FORCE training takes any high dimensional dynamical system and use this system as a computing reservoir. A series of readout weights (or decoders) for each neuron is learned online in a supervised fashion. Unlike the other two methods, one does not need the target behavior as a closed form differential equation to learn, merely a supervisor to provide an error signal to learn the decoders. The use of any high dimensional dynamical system as a reservoir makes the FORCE method applicable to many more network types while the use of an error signal expands the potential applications for any of these networks.

Unfortunately, FORCE training has only been directly implemented in networks of rate equations that phenomenologically represent how neurons firing rate varies smoothly in time, with little work done in spiking network implementations in the literature. A spiking implementation has been thought to be difficult so far and only two implementations have been proposed. DePasquale et. al., use a network of rate equations to train a spiking network [DePasquale et al., 2016]. The rate system is used to nonlinearly transform the target signal into a temporal basis set. This basis set is used as a target for the network of spiking neurons to learn. As this approach requires training the full connectivity matrix as opposed to just readout weights, it is computationally more difficult to apply. Additionally, Thalmier et. al., construct a spiking network of leaky integrate-and-fire neurons that implements the FORCE method to approximate linear and nonlinear dynamics. However, the network under consideration cannot perform nonlinear dynamics without nonlinear dendrites [Thalmeier et al., 2015]. Additionally, the network output and network type under consideration in [Thalmeier et al., 2015] have a very particular structure and architecture and it remains unclear whether FORCE training can generalize to other networks. Given the potential of the FORCE method in providing a novel and powerful top-down approach to network training, it is critical to see whether this technique can be directly applied to different networks consisting of different model neurons and what type of dynamics or functions these networks can learn.

We show that the FORCE method can be directly implemented in spiking networks and is robust against different implementations and neuron models. We then show that FORCE training can resolve many types of dynamical systems. Finally, we demonstrate that spiking neural networks can statistically classify inputs, and recall patterns after being trained with the FORCE method using a suitable teaching signal. These patterns include a sequence corresponding to the notes of a song, the spectrogram for the mating song of a zebra finch, and a short scene from a movie. Post-training inspection of the constituent neurons reveals behavior that is consistent with experimentally determined data such as the stereotypical decrease in voltage variance during input presentation [Churchland et al., 2010]. Thus, post-training analysis of the underlying neurons that form the network yields insight into how neuronal functioning relates to macroscopic and observable behaviors. We also found that the circuit constructed to recall our birdsong can display firing rate distributions that mirror experimentally obtained data from the primary nuclei responsible for songbird behavior, the High Vocal Center (HVC) and the robust nucleus of the arcopallium (RA). In the bird song circuit, area HVC provides a chain of synaptic activity to RA, which elicits bird song replay. We found that a similar chain of activity also facilitated learning and recall of a scene from a movie. Surprisingly, the network constructed to replay the movie displayed a prominent theta oscillation post-training in a proxy of its local field potential (LFP). By compressing, reversing, or attenuating the activity chain or clock, we could compress, reverse, or degrade replay of the movie respectively with corresponding changes to the LFP. This demonstrates a well defined, computational role of theta oscillations in learning and memory, as the temporal backbone on which memories are ordered, stored, and replayed [Buzsáki, 2002, Klimesch, 1999].

Results

We explored the potential of the FORCE method in training spiking neural networks to perform an arbitrary task. In these networks, the synaptic weight matrix is the sum of a set of static weights and a set of dynamically varying weights: $\boldsymbol{\omega} = \boldsymbol{\omega}^0 + \boldsymbol{\eta}\boldsymbol{\phi}^T$. The static weights, $\boldsymbol{\omega}^0$ are set to initialize the network into chaotic spiking [van Vreeswijk and Sompolinsky, 1996, Ostojic, 2014, Harish and Hansel, 2015]. The time dependent weights, $\boldsymbol{\phi}$, are learned using a supervised learning method called Recursive Least Mean Squares (RLMS) so that the squared error between the target dynamics (i.e. the task, $x(t)$) and the network dynamics ($\hat{x}(t)$) is minimized (Figure 1A) [Haykin, 2009, Bishop, 2006]. This method is successful when the network dynamics can mimic

the target dynamics even when RLMS is turned off.

FORCE Trained Networks Learns Dynamics Rapidly with Chaotic Rate Initialization

To demonstrate the basic principle of this method and to compare with our spiking network simulations, we applied FORCE training to a network of rate equations (Figure 1A) as in [Sussillo and Abbott, 2009] to learn a simple 5 Hz sinusoidal oscillator. The static weight matrix initializes high-dimensional chaotic dynamics onto the network of rate equations (Figure 1B). These dynamics form a suitable reservoir to allow the network to learn from a target signal quickly. As in the original formulation of the FORCE method, the rates are heterogeneous across the network and varied strongly in time (see Figure 2A in [Sussillo and Abbott, 2009]). The network is simulated first without RLMS for a short period of time to let the network settle in the chaotic attractor. Then, RLMS is activated (Figure 1B, C). After learning the appropriate weights ϕ_j (Figure 1 D), the network reproduces the oscillation without any guidance from the teacher signal, albeit with a slight frequency and amplitude error (Figure 1C).

To ascertain how these networks can learn to perform the target dynamics, we computed the eigenvalues of the resulting weight matrix before and after learning (Figure 1E). The eigenvalues lie within a circle whose radius can be predicted with results in random matrix theory [Rajan and Abbott, 2006]. As these eigenvalues are randomly distributed within a circle, there is no dominant eigenvalue (or singular value). This is a fundamentally different behavior then weight matrices generated by other top-down techniques such as the Neural Engineering Framework (NEF). NEF generated weight matrices are low rank for all-to-all coupled networks and thus have a small number of non-zero eigenvalues. The number of non-zero eigenvalues indicates the dimensionality of the target dynamics. FORCE trained weight matrices are always high rank and the number of non-zero eigenvalues is independent of the dimensionality of the target dynamics. However, if we scale up the effect of the target dynamics $\hat{x}(t)$ by increasing the magnitude of $\boldsymbol{\eta}\boldsymbol{\phi}^T$, then FORCE trained weight matrices can have dominant eigenvalues, but will still always be full rank. In these regimes, FORCE trained and NEF generated weight matrices are more comparable in structure and the dimensionality of the target dynamics is reflected by the number of dominant eigenvalues (see supplementary section S1 for the derivation of the NEF procedure for sparse networks, and supplementary Figures S1, and S2).

FORCE Trained Spiking Networks Learn Dynamics Rapidly with Chaotic Spiking Initialization

We sought to determine whether the FORCE method could train spiking neural networks given the impressive capabilities of this technique in training smooth systems of rate equations (Figure 1) previously demonstrated in [Sussillo and Abbott, 2009]. We implemented the FORCE method in different spiking neural networks of integrate-and-fire neurons in order to compare the robustness of the method across neuronal models. To demonstrate the chaotic nature of these networks, we deleted a single spike in the network [London et al., 2010]. After the deletion, the resulting spike trains immediately diverged across the network, which is a clear indication of chaotic behavior (Figure 2A). We then successfully trained networks of spiking theta neurons to mimic various types of oscillators such as sinusoids at different frequencies, sawtooth oscillations, Van der Pol oscillators, in addition to teaching signals with noise present (Figure 2B). With a 20 ms decay time constant, the product of sines oscillator with and without noise presented a challenge to the theta neuron to learn. However, it could easily be resolved for networks of theta neurons with larger decay time constants (see Supplementary Figure S3 and Supplementary Table S1). To ascertain whether or not refractory periods or spike-frequency adaptation would interfere in anyway with the learning process, we trained networks of Izhikevich neurons and leaky integrate-and-fire neurons with an absolute refractory period. Neither adaptation nor refractory periods interfered with the networks ability to learn from a supervisor (Figure 2C). All the oscillators in Figure 2B were trained successfully for both the Izhikevich and LIF neuron models (see Supplementary Figure S3 and Supplementary Table S1). Furthermore, FORCE training was robust to the possible types of initial chaotic network states (see Supplementary Figure S4) [Ostojic, 2014, Harish and Hansel, 2015].

As oscillators are very simple dynamical systems, we wanted to assess if this technique can train a spiking neural network to perform a much more complicated task, in this case reproducing the dynamics for a low-dimensional chaotic system. We trained a network of theta neurons using the Lorenz system as a teaching signal (Figure 2D). After 45 seconds of training, the network could reproduce the butterfly attractor and Lorenz-like trajectories after learning. As the teaching signal was much more complex then a mere oscillator, the training took longer and required more neurons yet was still successful.

In order to compare how the underlying structure of the network of spiking neurons with the network of rate neurons, we investigated how eigenvalues of the weight matrix varied in the spiking network before and

after learning. In the original rate based formulation, the final weight matrices could have either dominant eigenvalues or all the eigenvalues were uniformly distributed within a circle. The presence of dominant eigenvalues depended on the particular dynamics being learned, the relative strengths of the static component of the weight matrix and the learned component of the weight matrix. Here in the spiking network, we observe that in some cases, systems without dominant eigenvalues performed better then systems with dominant eigenvalues while in other cases the opposite was true (see Supplementary Material, Figure S2).

We wanted to determine if the FORCE training could be adapted to generate weight matrices that respect Dales law, the constraint that a neuron can only be excitatory or inhibitory, not both. Unlike previous work [DePasquale et al., 2016], we opted to enforce Dales law dynamically as the network was trained. We generated static, chaos inducing weight matrices that initially respected Dales law and enforced a boundary constraint on the dynamically varying weights $\boldsymbol{\eta}_i \boldsymbol{\phi}_j^T$. If the weight respected the excitatory/inhibitory identity of its presynaptic neuron, it was unaltered. If the weight violated this identity, it was set to 0 in the weight matrix. We tested this procedure with a network of Izhikevich neurons and successfully trained a sinusoidal oscillator (see Supplementary Material Section S2, Supplementary Figure S5). This implies that FORCE training can still successfully train a spiking neural network while respecting important biological constraints such as Dales law. For simplicity however, we train all remaining examples with unconstrained weight matrices.

For the dynamics under consideration, the Izhikevich model had the greatest accuracy and fastest training times. This is partially due to the fact that the Izhikevich neuron has the spike frequency adaptation variable (u_i) which operates on a longer time scale (i.e. 100 ms). The long time scale affords the reservoir a greater capability for memory, allowing it to learn longer signals. Additionally, the dimensionality of the reservoir is immediately increased due to having an extra adaptation variable for each neuron. Note that for all networks under consideration, we set the parameters such that the average firing rate was low (< 60 Hz), although the instantaneous firing rate of a single neuron may transiently increase above this. To summarize, for these three different neuron models, we have demonstrated that the FORCE method can be used to train a network to learn dynamics from a teaching signal. The teaching signal can be oscillatory, noisy, or even chaotic and the training can occur in a manner that respects Dales law.

The FORCE Method can Train Spiking Neural Networks to Classify Inputs

As populations of neurons encode more then just dynamical systems, we sought to determine what other potential behaviors or functions a simulated network could learn using FORCE training. For instance, human beings can naturally classify different objects as belonging to different categories or learn a particular sequence associated with the notes in a song. These tasks usually fall under the realm of artificial neural networks and the techniques used to train artificial neural networks are typically inapplicable in training an arbitrary spiking neural network [Bishop, 2006, Haykin, 2009].

To test whether a network of spiking neurons can function as a statistical classifier, we trained networks of Izhikevich neurons to classify inputs into two classes separated by either linear or nonlinear boundaries (Figure 3). The target dynamics are defined as an up-state when the input belongs to class 1 or a down state when the input belongs to class 2 (Figure 3A). The inputs are fed in sequentially to the network via feedforward weights. Nonlinear classification in a spiking neural network is similar in complexity to the XOR task studied in [DePasquale et al., 2016]. Each input is followed by a short break period where the network is trained to go to a rest state. Not only can the network perform the classification task, but it can generalize. Indeed, the network learns to classify new inputs after successive presentations of training data with FORCE training, albeit with some error. In both the linear and nonlinear cases, the classification test error rate was less then 7% on test data sets (Figure 3C).

As the majority of classifiers constructed with artificial neural networks are feedforward, we wanted to determine how the recurrent nature of the spiking network classifier influences the accuracy of the classification task. To investigate how the the network misclassifies potential inputs, we constructed peri-stimulus time histograms. The histograms were formed by repeatedly presenting inputs that were either far away from the true classification boundary or near the boundary (see supplementary material, Figure S6). Unlike a feedforward neural network, the recurrent neural network used here can classify the same input as either an up state or a down state. This depends on the initial condition of the network. Thus, some of the points near the boundary are misclassified in some cases, but correctly classified at later presentations (see supplementary material, Figure S6). Furthermore, when the network consistently classifies points as belonging to one class or another, the variance of the voltage traces decreases across the network (Figure 3D, Supplementary Figure S7). This result is consistent with the intracellular recordings from the primary visual cortex of cats [Churchland et al., 2010, Finn et al., 2007]. Indeed, the authors find a reduction in the voltage variance for stimuli that

are both preferred and nonpreferred by neurons. Here, we find that the consistency, and thus accuracy of the behavioral response in a classification task is mirrored by the voltage variance. When the network fails to consistently classify a point as an up or down state, the voltage variance during stimulus onset increases substantially.

The FORCE Method can Train Spiking Neural Networks to Repeat Complicated Temporal Signals

Neurons can encode complicated temporal sequences such as the mating songs that song birds learn, store, and replay [Hahnloser et al., 2002]. We wondered whether FORCE could train spiking networks to encode similar types of spatiotemporal sequences. This can essentially be formulated as a very long oscillation that is repeatedly presented to the network. The supervisors we considered were either a discrete sequence of up-states that the network had to learn or a high dimensional supervisor that corresponded to audio or visual data.

The first pattern we tested the network with was a sequence of up-states in a 5-dimensional teaching signal. These up-states appear in a specific order that corresponds to the first bar of Beethoven's Ode to Joy. The network can overall reproduce the sequence of notes in the first bar (see supplementary material, Figure 4A). Correct replay of the song corresponds to 82% of the signal. Note that the average firing rate of the network post-learning was 34 Hz, with variability in the neuronal responses from replay to replay, yet forming a stable peri-simulus time histogram (see supplementary material, Figure S8, S9, and S10).

While the network displayed some degree of error in replaying the song, the errors were not entirely random. The errors formed by this network are primarily located after the two non-unique E-note repetitions that occur in the first bar and the end of the third bar (Supplementary Figure S10) in addition to the non-unique ED sequences that occur at the end of the second bar and beginning of the fourth bar. The components of the sequence are referred to as high entropy in the songbird literature [Bouchard and Brainard, 2016]. It has been previously found that the post-stimulus activity of neurons in the song generating HVC area of birds is not predictive of upcoming syllables if the previously played stimulus was a song with high entropy regions. Here we find that the majority of the errors in recall occur during these high entropy transitions and can lead to incorrect replays of a wrong note sequence (see Supplementary Figures S8 and S10). Additionally, even for recalls classified as correct, the variance in the network output increases after these high entropy transitions (Supplementary Figure S10) indicating that the neuronal activity after these high entropy regions is not necessarily predictive of the behavioral output, as in the songbird data.

FORCE Trained Networks Can Reproduce Singing Behavior and Spiking Statistics of the Songbird Mating Song Circuit

We wondered if a more direct comparison to songbird singing behavior could be constructed using FORCE training. To that end, we analyzed and constructed a circuit that could reproduce a bird song recorded from an adult zebra finch. The singing behavior of these birds is owed to two primary nuclei: the High Vocal Center (HVC) and the robust nucleus of the arcopallium (RA). The HVC area consists of three primary classes of neurons: neurons that project to RA, neurons that project to the basal ganglia (referred to as area X projecting neurons) and inhibitory interneurons [Mooney, 2000, Kozhevnikov and Fee, 2007, Leonardo and Fee, 2005, Long et al., 2010, Hahnloser et al., 2002]. The RA projecting neurons in HVC form a chain of spiking activity and each RA projecting neuron fires only once at a specific time in the stereotypical bird song [Long et al., 2010, Hahnloser et al., 2002]. This chain of firing is transmitted to and activates the RA circuit. Each RA neuron bursts at precisely defined times during song replay and unlike HVC neurons, RA neurons burst multiple times during a song [Leonardo and Fee, 2005]. RA neurons then project to the lower motor nuclei that are responsible for generating the final bird song.

To simplify matters, we will focus on a network of RA neurons receiving a chain of inputs from area HVC (Figure 5A). The HVC chain of inputs are modeled directly as a series of successive up-states and are not FORCE trained. These up-states feed into a network of Izhikevich neurons that was successfully FORCE trained to reproduce the spectrogram of a recorded song from an adult zebra finch (Figure 5B, see Supplementary Video S1). Additionally, by varying the parameters governing the magnitude of the static weight matrix, the magnitude of the learned weight matrix, and some of the parameters for the Izhikevich model, the spiking statistics of RA neurons are easily reproduced both qualitatively and quantitatively (Figure 5D,E inset, see Figure 2B, 2C in [Leonardo and Fee, 2005], see Figure 1 in [Leonardo and Fee, 2005]) with RA neurons bursting regularly multiple times during song replay (Figure 5B,C). Thus, we have trained our model network to match the song generation behavior of RA with spiking behavior that is consistent with experimental data.

Given the correspondence between our model network and the RA nucleus [Leonardo and Fee, 2005], we wondered how our network would respond to potential pharmacological manipulations that could be applied to this area. In particular, we considered how upscaling or downscaling the amplitude of the excitatory synapses would effect the networks behaviors (Figure 5F). These manipulations and their resulting network behavior form experimentally testable predictions. They could be validated via the application of AMPA antagonists such as CNQX or GABA antagonists such as gabazine that can systematically alter the ratio between excitation and inhibition. We found that the network is robust towards downscaling excitatory synaptic weights, and still produced the song output, albeit at a lower amplitude even when the excitatory synapses were reduced by 90% of their amplitude (Figure 5F,G). However, upscaling excitatory synapses by as little as 15% drastically reduced song performance. The resulting song output had a very different spectral structure then the teaching signal as opposed to the downscaling excitation, which yielded a song that was highly correlated with the teaching signal. We were surprised at how robust the performance of the songbird network was given the high dimensionality and complicated structure of the output. We hypothesized that the performance of this network was strongly associated to the precise, clock-like inputs provided by HVC and that similar inputs could aid in the encoding and recall of other types of information. To test this hypothesis, we removed the HVC input pattern and found that the recall of the learned song was destroyed (not shown), consistent with experimental lesioning of this area in adult canaries [Nottebohm et al., 1976].

To further explore the benefits that these clock like signals might provide, we successfully FORCE trained a network of Izhikevich neurons to internally generate its own clock, while simultaneously also training it to reproduce the first bar of Ode to Joy. The entire supervisor consisted of the 5 notes used in the first bar of the song, in addition to 16 other components that correspond to a sequence of 16 up-states. The network learned both the clock signal and the song simultaneously, with less training time, and greater accuracy then without the clock signal (not shown). As the internal clock helped in learning and replaying the first bar of a song, we wondered if the clock signals could help in encoding longer signals. To test this hypothesis, we FORCE trained a network to learn the first four bars of Ode to Joy, corresponding to a 16 second song. The network successfully learned the song post-training, without any error in the sequence of notes in spite of the length and complexity of the sequence (see Supplementary Figure S11, Supplementary Video S2). Thus, internally generated clock signals for a network can make FORCE training faster, more accurate, and more robust to learning longer signals.

Theta Oscillations Help FORCE Trained Networks to Encode and Retrieve Episodic Memories

Given the improvements that a clock signal confers over the duration, training time, and accuracy, we wanted to know if these input signals would help populations of neurons to learn natural, high dimensional signals. To test this, we trained a network of Izhikevich neurons to learn a high dimensional teaching signal that corresponds to an 8 second clip from a movie (Figure 6). Each component of the teacher corresponds to the time evolution of a pixel in the movie scene. The clock signals were either generated by a separate network or were fed directly as an input into an encoding and replay network, as in the songbird example (Figure 6A, C). We were able to successfully train the network to replay the movie scene in both cases with a time averaged correlation coefficient of $r = 0.98$ (Figure 6B, see Supplementary Video S3) with occasional sharp and short dips in the correlation at specific times (not shown). Furthermore, we found that the clock inputs were necessary both for training and recall. The network could still recall the individual frames from the movie scene without the clock input, however the order of scenes was incorrect and appeared to be chaotic (see Supplementary Figure S12, Supplementary Video S3). Thus, clock inputs facilitate both learning and recall of high dimensional signals.

We were surprised to see that despite the complexity and high dimensionality of the encoding signal, the histogram of spike times for the network displayed a strong theta-modulation. This histogram can be thought of as a proxy for the local field potential of a network (LFP). FORCE trained networks provide a direct link between how manipulations of the circuit affect both the behavioral output (movie replay) and other measurable quantities such as our LFP proxy. Given these capabilities, we wanted to determine how manipulating either the clock network, or the recall would effect both recall performance and the behavior of the LFP proxy. Unsurprisingly, reducing the clock amplitude yields a steady decrease in recall performance. Initially, this is mirrored through a decrease in the amplitude of the theta oscillations in the LFP proxy (See Supplementary Figure S12). Surprisingly however if we completely remove the clock inputs then the LFP proxy exhibits a strong delta wave oscillation ($\approx 2Hz$) (Figure 6 G). The emergence of these delta oscillations correspond to a sharp decline in recall performance. Note that the network still stores the individual scenes in this movie, but can no longer replay them in the correct order without the clock input (See Supplementary Video S4). Thus,

the theta oscillation in the local field potential is critical for encoding and replay of the movie scene, which can serve as a proxy for an episodic memory.

We wanted to determine how lesioning neurons in the recall network would affect both the LFP, and recall performance. We found that the recall performance decreased approximately linearly with the proportion of neurons that were lesioned, with the amplitude of the LFP also decreasing (Figure 6 F, G). This mirrors experimental results that showed that theta power was predictive of correct recall [Lisman and Jensen, 2013]. The clock network however was much more sensitive to neural lesioning. We found lesioning 10% of a random selection of neurons in the clock network was sufficient to stop the clock output. Thus, the clock network is the critical element in this circuit that can drive accurate replay and is the most sensitive to damaging perturbations such as neural loss.

It has been speculated that accelerated replay of a scene or sequence may be important for memory consolidation [Euston et al., 2007], thus we wondered if pre-trained networks with a specific clock input could replay the scene in accelerated time. To test this hypothesis, we started with a network that was trained with the default clock with period 8 seconds and compressed the clock in time (Figure 6D, Supplementary Video S5). The network was able to replay the movie in compressed time with a correlation of > 0.8 up to a compression factor of $16\times$ with accuracy sharply dropping for further compression. The effect of time compression on our LFP proxy was to introduce higher frequency oscillations (see Supplementary Figure S13, Figure 6G). The frequency of these oscillations scaled linearly with the amount of compression. However, with increasing compression frequency, large waves of synchronized activity also emerged in the LFP proxy. These bear a similarity to hippocampal sharp waves, which also occur with compressed firing sequences of place cells for example [Buzsáki, 2015]. We should note however that the stereotypical ripple complex that simultaneously occurs with sharp waves is not present in our results. This is likely due to a lack of a specific excitatory and inhibitory subpopulations. As reversed sequence replay is also speculated to be important for learning, we wondered if we could also reverse the replay of the movie [Diba and Buzsáki, 2007]. This was successfully achieved by reversing the order in which the clock components were presented to the network with an accuracy of $r = 0.90$ (see Supplementary Video S6). Thus, the network can generalize to robustly compress in time or reverse replay, despite not being trained to do these tasks.

Compression of a task dependent sequence of firing has been experimentally found in numerous sources ([Euston et al., 2007, Mello et al., 2015, Nádasdy et al., 1999]). For example, in [Euston et al., 2007], the authors found that a recorded sequence of neuronal cross correlations in rats elicited during a spatial sequence task reappeared in compressed time during sleep. The compression ratios between 5.4-8.1 and compressed sequences that were originally ≈ 10 seconds down to ≈ 1.5 seconds, which is similar to the compression ratios we found for our networks without incurring appreciable error in replay. Time compression and dilation has also been experimentally found in the striatum [Motanis and Buonomano, 2015, Mello et al., 2015]. Here, the authors found that the neuronal firing sequences for striatal neurons were directly compressed in time in response to a fixed interval task with correspondingly shorter fixed intervals ([Mello et al., 2015]). Indeed, we also found that accelerated replay of our movie forced the voltage traces to have compressed patterns of firing in time (see Supplementary Figure S13).

Discussion

We have shown that FORCE training can take initially chaotic networks of spiking neurons and use them to perform a wide variety of dynamics and tasks that mimic the natural tasks and functions demonstrated by populations of neurons. For example, these networks were trained to learn low-dimensional dynamical systems such as oscillators. Low-dimensional oscillatory systems are at the heart of generating not only rhythmic motion, but also non-rhythmic motion such as the reaching tasks observed in motor cortex [Churchland et al., 2012]. Additionally, we showed that we could train neural networks to display behaviors that are beyond simple dynamical systems by altering the teaching signal used to train the network. For example, we trained a statistical classifier with a network of Izhikevich neurons that could discriminate populations that were either linearly or nonlinearly separable. Extending the notion of an oscillator even further allowed us to store a complicated sequence of events in the form of notes of a song, reproduce the singing behavior of songbirds, and encode and replay a movie scene.

FORCE Trained Networks and Songbird Behaviors

Our FORCE training procedure is reminiscent of how songbirds learn their stereotypical learned songs [Konishi, 1985, Hahnloser et al., 2002]. Songbirds are typically presented with a species specific song or repertoire of

songs from their parents or other members of their species when they are young. These juvenile birds internalize the original template song and subsequently use it as an error signal for their own vocalization, similar to how we have trained our network of spiking neurons to learn the first bar of Ode to Joy [Konishi, 1985].

We reproduced the singing behavior of songbirds with a FORCE trained network [Konishi, 1985, Mooney, 2000, Bouchard and Brainard, 2016, Kozhevnikov and Fee, 2007, Leonardo and Fee, 2005, Long et al., 2010, Hahnloser et al., 2002, Nottebohm et al., 1976]. We used a sequence of up-states as an input to a network of spiking neurons as our model of the chain of activity that neurons in RA receive from projection neurons in area HVC. Both the spiking statistics of area RA and the song spectrogram were accurately reproduced after FORCE training. Furthermore, we demonstrated that altering the balance between excitation and inhibition post training might affects the singing behavior. A shift to excess excitation alters the spectrogram instead of merely amplifying it while a shift to excess inhibition reduces amplitude of all frequencies in the spectrogram.

Furthermore, we demonstrated a strong link between how the songbird circuit functions and episodic memories are learned and replayed using a theta oscillation. In both networks, a chain of activity serves as a clock to encode time and this clock elicits correct ordered replay in down-stream populations of neurons. The nature of this chain can differ from neural circuit to neural circuit. For example, there is an explicit firing chain from HVC RA projection neurons which results in ordered bursting in the RA network. Our episodic and Ode to Joy recall networks however incorporated population chains that also operate on a slower time scale then the HVC RA projection chain. In these chains, each neuron can fire multiple times in the overall sequence, yet still in a specific order with regards to other neurons in the network. The Ode to Joy network simultaneously encoded this population chain with the actual sequence of notes while there was a separate population chain or clock network in the movie recall example. The resulting spike time histograms across the networks display a clear theta oscillation. It is worth noting that the circuit responsible for the singing behavior in songbirds is also present in parrots and humminbirds, and is thought to have evolved independently in all three species [Jarvis et al., 2000].

Episodic Memories can be Encoded and Replayed with FORCE Training Using Theta Oscillations

Inspired by the clock like input pattern that song birds use for learning and replay [Long et al., 2010, Hahnloser et al., 2002] we used similar signals to encode a longer and more complex sequence of notes in addition to a scene from a movie. These clock signals could be internally generated and learned by the network as in the longer Ode to Joy example, or fed from an external source as in the movie replay example. We found that these clock signals made both FORCE training faster and the subsequent recall more accurate. Furthermore, by manipulating the clock frequency we found that we could speed up or reverse movie replay in a robust fashion. We found that compressing replay resulted in higher frequency oscillations in our LFP proxy. Attenuating the clock signal decreases recall performance while transitioning our LFP proxy from a theta oscillation to a delta oscillation. Finally, replay of the movie was robust to lesioning neurons in the recall network.

While the theta oscillation is strongly associated to memory, it computational role is not fully understood, with many theories proposed [Buzsáki, 2002, Klimesch, 1999, Buzsáki, 2005, Lengyel et al., 2005]. For example, it is well known that the theta oscillation promotes the expression of long term potentiation (LTP), with increasing theta power increasing the strength of the induced LTP [Klimesch, 1999] and presumably, the strength of memory formation. Additionally, the theta oscillation has been proposed to serve as a clock for memory formation [Lubenov and Siapas, 2009, Buzsáki, 2002]. Here, we show a concrete example that very complicated and high dimensional memories can be bound to an underlying theta oscillation. The oscillation functions as a clock or temporal backbone in a very explicit mechanism for storage and replay. The underlying clock can be sped up, or even reversed for longer term memory consolidation. Additionally, we found that blocking this clock input severely disrupted learning and disrupted the underlying theta oscillation. Blocking the hippocampal theta oscillation pharmacologically ([Givens and Olton, 1995]) or optogenetically ([Boyce et al., 2016]) has also been found to disrupt learning. In [Boyce et al., 2016] for example, the hippocampal theta oscillation was blocked by optogenetically silencing a set of inhibitory neurons in the medial septum only during REM sleep. The result of this silencing is abolishing theta oscillations across the hippocampus. The REM specific blocking of theta oscillations damaged memory consolidation for novel object recognition and fear conditioned contextual memory. This demonstrates the necessity of the theta oscillation for consolidating a memory trace, mirroring our results where successive presentations of the movie signal were necessary to consolidate our movie into memory during FORCE training.

The potential existence for a subset of neurons whose sequential firing can serve as a backbone has recently been observed in a subset of hippocampal CA1 pyramidal neurons [Grosmark and Buzsáki, 2016]. These neurons had low spatial specificity but still took part in sequential firing during sharp-wave ripples in the

hippocampus. They displayed very limited changes in their sequential firing after sleep while a remaining plastic set of neurons displayed a higher resulting place specificity after sleep. Thus, [Grosmark and Buzsáki, 2016] conclude that sequence replay during sharpwave ripples occurs due to a subset of plastic neurons being added to a rigid backbone sequence. While the exact details may differ between our episodic memory example, the core idea of using sequential firing of a population of neurons as a backbone for the encoding of new information is precisely the mechanism found in [Grosmark and Buzsáki, 2016].

FORCE Training and Other Top Down Approaches

FORCE training is a very powerful tool that allows one to use any sufficiently complicated dynamical system as a basis for universal computation. Given the potential power of this technique in controlling and harnessing the omnipresent complexity of a large spiking neural network, it is naturally surprising that there is very limited work done in spiking network implementations. The primary difficulty in implementing the technique appears to be controlling the orders of magnitude between the chaos inducing weight matrix and the feedback weight matrix. If the chaotic weight matrix is too large in magnitude (via the G parameter), the chaos can no longer be controlled by the feedback weight matrix [Sussillo and Abbott, 2009]. However, if the chaos inducing matrix is too small, then the chaotic system no longer functions as a suitable reservoir. Thus, the suitable values of G are at the so-called edge of chaos, near the regions where chaos can no longer be tamed. Furthermore, while we succeeded in implementing the technique in other neuron types, the Izhikevich model was the most accurate in terms of learning arbitrary tasks or dynamics. This is due to the presence of spike frequency adaptation variables that operate on a much slower time scale then the neuronal equations. These slowly changing adaptation currents increase the dimensionality of the reservoir while providing it with an important source of memory through their long time scale dynamics. There may be other biologically relevant forces that can increase the capacity of the network to act as a reservoir through longer time scale dynamics, such as synaptic depression and NMDA mediated currents for example.

Although FORCE trained networks have dynamics that are starting to resemble those of populations of neurons, at present all top-down procedures used to construct any functional spiking neural network need further work to become biologically plausible learning rules [Sussillo and Abbott, 2009, Boerlin et al., 2013, Eliasmith et al., 2012]. For example, FORCE trained networks require non-local information in the form of the correlation matrix $\boldsymbol{P}(t)$. However, we should not dismiss the final weight matrices generated by these techniques as biologically implausible simply because the techniques are themselves biologically implausible. More work should be done in implementing either FORCE, NEF, or spike-based coding networks using a biologically plausible learning mechanism based on synaptic plasticity or homeostasis [Bi and Poo, 1998, Pfister and Gerstner, 2006, Clopath et al., 2010, Graupner and Brunel, 2012, Babadi and Abbott, 2016, Vogels et al., 2011]. This has been resolved for spike-based coding networks and linear dynamical systems for example [Bourdoukan and Denève, 2015]

The versatility of FORCE learning is due to the fact that the technique is an implementation of reservoir computing from the machine learning field [Lukoševičius et al., 2012, LukošEvičIus and Jaeger, 2009, Schrauwen et al., 2007]. In reservoir computing, a recurrently coupled dynamical reservoir takes some dynamically varying input and transforms it into a much higher dimensional dynamical system. As these techniques make minimal assumptions about the reservoir, they can be applied to a wide variety of spiking neural networks. Reservoir techniques work by finding a set of readout weights on the reservoir to yield the target dynamics. This approach vastly simplifies the training of recurrently coupled spiking networks as only the readout weights have to be learned, as opposed to the recurrent weights. In the FORCE method and other implementations of reservoir computing, the network approximant is fed back into the network. The general approach to reservoir computing has many origins [Dominey, 1995, Maass et al., 2002, Jaeger, 2001]. What constitutes a reservoir and more importantly what constitutes a good reservoir is currently an active field of research [Schliebs et al., 2011, Ozturk and Principe, 2005, Maass, 2010, Maass et al., 2004, Wojcik and Kaminski, 2004]. The FORCE method uses a balanced, chaotic network as a reservoir, with similar ideas for general chaotic reservoirs appearing previously in [Ozturk and Principe, 2005], in addition to chaotic spiking networks of excitatory and inhibitory neurons in [Lourenço, 2006]. The earliest implementation of a spiking network reservoir appears to be in [Buonomano and Merzenich, 1995].

Aside from the original rate formulation in [Sussillo and Abbott, 2009], FORCE trained rate equations have been recently applied to analyzing and reproducing experimental data. For example, in [Rajan et al., 2016], the authors used a variant of FORCE training (referred to as Partial In-Network Training, PINning) to train a rate network to reproduce a temporal sequence of activity from mouse calcium imaging data. PINning uses minimal changes from a balanced weight matrix architecture to form neuronal sequences. Additionally,

in [Sussillo et al., 2015], the authors used FORCE training on a rate network to reproduce the firing rates and population level responses of monkey grasping originally obtained in [Churchland et al., 2012]. The authors found that the resulting networks learned low dimensional oscillators, which were also present in the jPCA (a variant of principal component analysis) reduced data from [Churchland et al., 2012]. In [Li et al., 2016], the authors combine experimental manipulations with FORCE trained networks to demonstrate that preparatory activity prior to motor behavior is resistant to unilateral perturbations both experimentally, and in their FORCE trained rate models. Finally, in [Laje and Buonomano, 2013], the authors train a network of rate neurons to encode time on the scale of seconds. This network is subsequently used to learn different spatio-temporal tasks, such as a cursive writing task and were able to account for psychophysical results such as Weber's law, where the variance of a response scales like the square of the time since the start of the response. In all cases, FORCE trained rate networks were able to account for and predict experimental findings. Thus, FORCE trained spiking networks can prove to be invaluable for generating new predictions using voltage traces, spike times, and neuronal parameters.

The FORCE method differs from the other top-down approaches for generating neural networks by treating the optimization procedure differently. In both the Neural Engineering Framework (NEF) and spike-based coding approaches, the networks are constructed by solving different optimization problems. For both techniques, the weights are immediately solved for in a single step. This is prior to simulating the network rather then iteratively in an online fashion like the FORCE method. However, there is some promising recent work for spike based coding in learning linear dynamical systems iteratively in a biologically plausible manner [Bourdoukan and Denève, 2015]. The NEF approach converts the discrete spiking into rates by using the frequency current ($f(I)$) curve for each neuron in the network. The weights for each individual $f(I)$ curve can be solved for immediately so that the network obeys a closed-form differential equation that the user selects. Spike-based networks assume that the error between the intended dynamics and the networks approximant is encoded in the membrane potential of each neuron in the network. The procedure involves solving a time varying optimization problem that determines not only the weight matrix, but which neuron model to use [Boerlin et al., 2013, Schwemmer et al., 2015]. One of the surprising results is that even though the spike-based coding networks employ a temporal optimization, the solution obtained is a static set of connection weights and it does not require simulating the network first.

While the NEF and spike based coding approaches provide immediate weight matrix solutions, both techniques are difficult to generalize to other types of networks or other types of tasks. For the NEF approach, one has to know how neurons convert currents into spike firing rates via the $f(I)$ curve and these curves must be static to perform the optimization. Furthermore, one has to know the dynamics of the synaptic connection type to apply the optimization (see Supplementary Material Section S1). This limits the applicability of NEF generated networks to specific neuron types and specific synaptic models. For the spike-based coding approach, the entire network is determined by the task or dynamics for the network to perform [Boerlin et al., 2013]. While there is some promising recent work in applying spike-based coding approaches to more biologically plausible neurons [Schwemmer et al., 2015], the networks under consideration have not been extended to other synapse types or include other potential biologically relevant forces. Finally, both the NEF and spike based coding approach require a system of closed form differential equations to determine the static weight matrix that yields the target dynamics. This limits the types of tasks these networks can perform.

With an underlying spiking neural network that can reproduce realistic dynamics and behaviors, we can analyze the networks components for comparisons with recorded data. For example, our classifier displayed the characteristic decrease in variance upon input presentation ([Churchland et al., 2010]), the output of our song network was consistent with spiking data from area RA ([Leonardo and Fee, 2005]) while our movie replay network had compression factors consistent with experimentally verified data ([Euston et al., 2007]) and reproduced temporally compressed firing sequences ([Mello et al., 2015]). More importantly however, we are able to generate powerful new predictions such as the role of the theta oscillation as a clock or temporal backbone upon which episodic memories are encoded.

Methods

Rate Equations

The network of rate equations is given by the following:

$$\tau_s \dot{s}_i = -s_i + G\sum_{j=1}^{N} \omega_{ij}^0 r_j + Q\eta_i \hat{x} \tag{1}$$

$$r_j = \begin{cases} F\sqrt{s_j} & s_j \geq 0 \\ 0 & s_j < 0 \end{cases} \tag{2}$$

$$\hat{x}(t) = \sum_{j=1}^{N} \phi_j r_j(t) \tag{3}$$

where ω_{ij}^0 is the sparse and static weight matrix that induces chaos in the network by having ω_{ij}^0 drawn from a normal distribution with mean 0 and variance $(Np)^{-1}$, where p is the degree of sparsity in the network. The variables s can be thought of as neuronal currents with a postsynaptic filtering time constant of τ_s. The encoding variables η_i are drawn from a uniform distribution over $[-1, 1]$ and set the neurons encoding preference over the phase space of $x(t)$, the target dynamics. The quantities G and Q are constants that scale the static weight matrix and feedback term, respectively. The firing rates are given by r_j and correspond to the Type-I normal form for firing. The variable F scales the firing rates. The decoders, ϕ_j are determined by the Recursive Least Mean Squares (RLMS) technique, iteratively [Haykin, 2009, Bishop, 2006]. RLMS is described in greater detail in the next section.

Integrate-and-Fire Networks

Our networks consist of coupled integrate-and-fire neurons, that are one of the following three forms:

$$\dot{\theta}_i = (1 - \cos(\theta_i)) + \pi^2(1 + \cos(\theta_i))(I) \quad \text{(Theta neuron)} \tag{4}$$

$$\dot{v}_i = -\frac{v_i}{\tau_m} + I \quad \text{(LIF Neuron)} \tag{5}$$

$$C\dot{v}_i = k(v_i - v_r)(v_i - v_t) - u_i + I \quad \text{(Izhikevich Neuron)} \tag{6}$$

$$\dot{u}_i = a(b(v_i - v_r) - u_i) \tag{7}$$

The currents, I are given by $I = I_{Bias} + s_i$ where I_{Bias} is a constant background current set near or at the rheobase (threshold to spiking) value. The currents are dimensionless in the case of the theta neuron while dimensional for the LIF and Izhikevich models. Note that we have absorbed a unit resistance into the current for the LIF model. The quantities θ_i, or v_i are the voltage variables while u_i is the adaptation current for the Izhikevich model. These voltage variables are instantly reset to a reset potential (v_{reset}) when a neurons membrane potential reaches a threshold (v_{thr},LIF) or voltage peak (v_{peak}, theta model, Izhikevich model). The parameters for the neuron models can be found in Table 1. The LIF model has a refractory time period, τ_{ref} where the neuron cannot spike. The adaptation current, u_i for the Izhikevich model increases by a discrete amount d every time neuron i fires a spike and serves to slow down the firing rate. The membrane time constant for the LIF model is given by τ_m. The variables for the Izhikevich model include C (membrane capacitance), v_r (resting membrane potential), v_t (membrane threshold potential), a (reciprocal of the adaptation time constant), k (controls action potential half-width) and b (controls resonance properties of the model).

The spikes are filtered by specific synapse types. For example, the single exponential synaptic filter is given by the following:

$$\dot{r}_j = -\frac{r_j}{\tau_s} + \frac{1}{\tau_s}\sum_{t_{jk}<t} \delta(t - t_{jk}) \tag{8}$$

where τ_s is the synaptic time constant for the filter and t_{jk} is the kth spike fired by the jth neuron. The double exponential filter is given by

$$\dot{r}_j = -\frac{r_j}{\tau_d} + h_j \tag{9}$$

$$\dot{h}_j = -\frac{h_j}{\tau_r} + \frac{1}{\tau_r \tau_d}\sum_{t_{jk}<t} \delta(t - t_{jk}) \tag{10}$$

where τ_r is the synaptic rise time, and τ_d is the synaptic decay time. The filters can be of the simple exponential type, double exponential, or alpha synapse type ($\tau_r = \tau_d$). However, we will primarily consider the double exponential filter with a rise time of 2 ms and a decay time of 20 ms.

The synaptic currents, $s_i(t)$, are given by the equation

$$s_i = \sum_{j=1}^{N} \omega_{ij} r_j \tag{11}$$

The matrix ω_{ij} is the matrix of synaptic connections and controls the magnitude of the postsynaptic currents arriving at neuron i from neuron j.

The primary goal of the network is to approximate the dynamics of an m-dimensional teaching signal, $\boldsymbol{x}(t)$, with the following approximant:

$$\hat{\boldsymbol{x}}(t) = \sum_{j=1}^{N} \boldsymbol{\phi}_j r_j \tag{12}$$

where $\boldsymbol{\phi}_j$ is a quantity that is commonly referred to as the linear decoder for the firing rate. The FORCE method decomposes the weight matrix into a static component, and a dynamically varying decoder $\boldsymbol{\phi}$:

$$\omega_{ij} = G\omega_{ij}^0 + Q\boldsymbol{\eta}_i \cdot \boldsymbol{\phi}_j^T \tag{13}$$

The matrix sparse ω_{ij}^0 is the static weight matrix that induces chaos and is drawn from a normal distribution with mean 0, and variance $(Np)^{-1}$, where p is the sparsity of the matrix. Additionally, the sample mean for these weight matrices was explicitly set to 0 for the LIF and theta neuron models to counterbalance the finite size effect. This was not necessary with the Izhikevich model. The variable G controls the chaotic behavior and its value varies from neuron model to neuron model. The encoding variables, $\boldsymbol{\eta}_i$ are drawn randomly and uniformly from $[-1, 1]^k$, where k is the dimensionality of the teaching signal. The variable Q is increased to better allow the desired recurrent dynamics to tame the chaos present in the reservoir. The weights are determined dynamically through time by minimizing the squared error between the approximant and intended dynamics, $\boldsymbol{e}(t) = \hat{\boldsymbol{x}}(t) - \boldsymbol{x}(t)$. The Recursive Least Mean Squares (RLMS) technique updates the decoders accordingly:

$$\boldsymbol{\phi}(t) = \boldsymbol{\phi}(t - \Delta t) - \boldsymbol{e}(t)\boldsymbol{P}(t)\boldsymbol{r}(t) \tag{14}$$

$$\boldsymbol{P}(t) = \boldsymbol{P}(t - \Delta t) - \frac{\boldsymbol{P}(t - \Delta t)\boldsymbol{r}(t)\boldsymbol{r}(t)^T\boldsymbol{P}(t - \Delta t)}{1 + \boldsymbol{r}(t)^T\boldsymbol{P}(t - \Delta t)\boldsymbol{r}(t)} \tag{15}$$

where $\boldsymbol{P}(t)$ is the network estimate for the inverse of the correlation matrix:

$$\boldsymbol{P}(t)^{-1} = \int_0^t \boldsymbol{r}(t)\boldsymbol{r}(t)^T \, dt + \lambda I_N \tag{16}$$

and λ is a regularization parameter [Haykin, 2009] and I_N is an $N \times N$ identity matrix. The RLMS technique can easily be seen to minimize the squared error between the target dynamics and the network approximant:

$$C = \int_0^T (\hat{x}(t) - x(t))^2 \, dt + \lambda \boldsymbol{\phi}^T \boldsymbol{\phi} \tag{17}$$

and uses network generated correlations optimally by using the correlation matrix $P(t)$. The network is initialized with $\boldsymbol{\phi}(0) = \boldsymbol{0}$, $\boldsymbol{P}(0) = I_N\lambda^{-1}$, where I_N is an N-dimensional identity matrix.

Further Methods for Each Simulation

The FORCE Method with Rate Networks (Figure 1)

A network of rate equations is used to learn the dynamics of a teaching signal, $x(t)$, a 5 Hz sinusoidal oscillation. The network consists of 1000 rate neurons with $F = 10$, $\tau_s = 10$ ms, $G = 1$, and $Q = 1.5$, $p = 0.1$. The network is run for 1 second initially, with RLMS run for the next 4 seconds, and shut off at the 5 second mark. The network is subsequently run for another 5 seconds to ensure learning occurs. The smooth firing rates of the network were less then 30 Hz. The RLMS parameters were taken to be $\Delta t = 2$ ms and $\lambda^{-1} = 0.5$ ms.

Using Spiking Neural Networks to Mimic Dynamics Using FORCE training (Figure 2)

Figure 2B

Networks of theta neurons were trained to reproduce different oscillators. The parameters were: $N = 2000$, $G = 10$ (sinusoid, sawtooth), $G = 15$ (Van der Pol), $G = 25$ (product of sinusoids), $G = 15$ (product of sinusoids with noise), $Q = 10^4$, $\Delta t = 0.5$ ms, with an integration time step of 0.01 ms. The resulting average firing rates for the network were 26.1 Hz (sinusoid) 29.0 Hz (sawtooth), 15.0 Hz (Van der Pol relaxation), 18.8 Hz (Van der Pol Harmonic), 23.0 Hz (product of sinusoids), 21.1 Hz (product of sines with noise). The Van der Pol oscillator is given by: $\ddot{x} = \mu(1 - x^2)\dot{x} - x$ for $\mu = 0.3$ (harmonic) and $\mu = 5$ (relaxation) and was rescaled in space to lie within $[-1, 1]^2$ and sped up in time by a factor of 20. The networks were initialized with the static matrix with 10% connectivity that induces chaos for 5 seconds, underwent another 5 seconds of FORCE training, after which RLMS was deactivated for the remainder 5 seconds of simulation time for all teaching signals except signals that were the product of sinusoids. These signals required 50 seconds of training time and were tested for a longer duration post training to determine convergence. The λ^{-1} parameter used for RLMS was taken to be the integration time step, 0.01 ms for all implementations of the theta model in Figure 2. The integration time step used in all cases of the theta model was 0.01 ms.

Figure 2C

The theta neuron parameters were as in Figure 2B for the 5Hz sinusoidal oscillator. The LIF network consisted of 2000 neurons with $G = 0.04$, $Q = 10$, with 10% connectivity in the static weight matrix and $\Delta t = 2.5$ ms with an integration time step of 0.5 ms. The average firing rate was 22.9 Hz. The Izhikevich network consisted of 2000 neurons with $G = 5 * 10^3$, $Q = 5 * 10^3$, with 10% connectivity in the static weight matrix and $\Delta t = 0.8$ ms and an average firing rate of 36.7 Hz. The λ^{-1} parameter used for RLMS was taken to be a small fraction of the integration time step, $\lambda^{-1} = 0.0025$ ms for Figure 2 while the Izhikevich model had $\lambda^{-1} = 1$ ms and an integration time step of 0.04 ms.

Figure 2D

The Lorenz system with the parameter $\rho = 28$, $\sigma = 10$, and $B = 8/3$ was used to train a network of 5000 theta neurons and is given by the equations:

$$\begin{aligned} \dot{x} &= \sigma(y - x) \\ \dot{y} &= x(\rho - z) - y \\ \dot{z} &= xy - Bz. \end{aligned}$$

The connectivity parameters were $G = 14$, and $Q = 10^4$. The average firing rate was 22.21 Hz. The network was trained for 45 seconds, and reproduced Lorenz-like trajectories and the attractor for the remaining 50 seconds.

Using Spiking Neural Networks for Classification with FORCE training (Figure 3)

Both the linear and nonlinear classification problems were resolved with a network fo 2000 Izhikevich neurons with the same synaptic and neuronal parameters as before. The remaining parameters were $G = 6 * 10^3$, $Q = 5 * 10^3$, and $\Delta t = 0.8$ ms and an integration time step of 0.04 ms. The teaching signal was constructed using the positive component of a sinusoid with a frequency of 2 Hz. The input data was uniformly distributed inside $[0, 1]^2$, with the classification boundary being $y = x$ in the linear case, and $y = \sin(2\pi x)$ in the nonlinear case. The inputs were multiplied by an $N \times 2$ input weight matrix W^{in} to yield the input current where each row of W^{in} is drawn uniformly from the circle with radius 500. The average rate for these networks was 12.5 Hz. $\lambda^{-1} = 30$ ms was used for RLMS training. A total of 800 seconds of training time was used corresponding to 1600 inputs.

Using Spiking Neural Networks for Pattern Storage and Recall with the FORCE Method (Figure 4)

The teaching signal was constructed using the positive component for a sinusoid with a frequency of 2 Hz for quarter notes and 1 Hz for half notes. Each upstate corresponds to the presence of a note and they can also be used as the amplitude/envelope of the harmonics corresponding to each note to generate an audio-signal from the network (see Supplementary Material for Audio File). The network consisted of 5000 Izhikevich neurons

with identical parameters as before, $G = 1 * 10^4$, $Q = 4 * 10^3$, and $\Delta t = 4$ ms and an integration time step of 0.04 ms. The teaching signal was continually fed into the network for 900 seconds, corresponding to 225 repetitions of the signal. After training, the network was simulated with RLMS off for a period of 1000 seconds. Within the 1000 seconds, correct replay of the song corresponded to 82% of the duration with 205 distinct, correct replays of the song within the signal. Correct replays were automatically classified by constructing a moving average error function:

$$E(t) = \sum_{i=1}^{5} \int_{t}^{t+T} (\hat{x}_i(t') - x_i(t'))^2 \, dt' \tag{18}$$

where $T = 4$ seconds is the length of the song. As we vary t, the teaching signal $x(t')$ comes into alignment with correct playbacks in the network output, $\hat{x}(t')$. This reduces the error creating local minima of $E(t)$ that correspond to correct replays at times $t = t^*$. These are automatically classified as correct if $E(t^*)$ is under a critical value. The average firing rate for the network was 34 Hz. $\lambda^{-1} = 2$ ms was used for RLMS training.

Using Spiking Neural Networks for Pattern Storage and Recall with the FORCE Method: Song Bird Example (Figure 5)

The teaching signal consisted of the spectrogram for a 5 second long recording of the singing behavior of a male zebra finch that was generated with a 22.7 ms window. This data was obtained from the CRCNS data repository [Crandall and Nick, 2014]. The RA network consisted of 1000 neurons that was FORCE trained for 50 seconds with $G = 1.3 * 10^4$, $Q = 1 * 10^3$, $\lambda^{-1} = 2$ ms, $\Delta t = 4$ ms and an integration time step of 0.04 ms. The average firing rate for the network was 52.93 Hz. The HVC input was modeled as a sequence of 500 upstates that were constructed from the positive component of a sinusoidal oscillation with a period of 20 ms. The up-states fired in a chain starting with the first component. This chain of upstates was multiplied by a static, $N \times 500$ feedforward weight matrix, W^{in} where each weight was drawn uniformly from the $[-8 * 10^3, 8 * 10^3]$. The network was tested for another 50 s after RLMS was turned off to test if learning was successful. We note that for this songbird example, the network could also learn the output with only the static, chaos inducing weight matrix and the feedforward clock inputs ($Q = 0$). However when $Q = 0$, the network behavior is similar to the more classic liquid state regime.

Using Spiking Neural Networks for Pattern Storage and Recall with the FORCE Method: Movie Replay (Figure 6)

The teaching signal consisted of a movie clip from the movie Predator with an 8 second duration. The frames were smoothed and interpolated so that RLMS could be applied at any time point, instead of just the fixed times corresponding to the native frames per second of the original movie clip. The movie was downsampled to 1920 pixels (30×64) which formed the supervisor signal. For both implementations (autonomous and non-autonomous recall), the recall network consisted of 1000 neurons with $G = 5 * 10^3$, $Q = 4 * 10^2$, $\lambda^{-1} = 2$ ms, $\Delta t = 4$ ms and an integration time step of 0.04 ms. For the non-autonomous recall, a clock signal was generated with 32 up-states that were 250 ms in duration. The chain of upstates was multiplied by a static, $N \times 32$ feedforward weight matrix, W^{in} where each weight was drawn uniformly from the $[-4 * 10^3, 4 * 10^3]$ and 74 seconds of FORCE training was used. The network was simulated for another 370 s after RLMS was turned off to test if learning was successful. For autonomous recall, a separate clock network was trained with 2000 neurons. The supervisor for this network consisted of a clock signal with 64 up-states that were 125 ms in duration. The chain of upstates was fed into an identical recall network as in the non-autonomous recall task. The recall network was trained simultaneously to the clock network and had $G = 1 * 10^4$, $Q = 4 * 10^3$. The clock network output was fed into the recall network in the same way as the non-autonomous recall task. RLMS was applied to train the clock and recall network with identical parameters as in the non-autonomous recall, only with a longer training time (290 s). We note that for this movie example, the network could also learn the output with only the static, chaos inducing weight matrix and the feedforward clock inputs ($Q = 0$). However when $Q = 0$, the network behavior is similar to the more classic liquid state regime.

Neuron Model	Parameter	Value
Izhikevich	C	250 μ F
	v_r	-60 mV
	v_t	-20 mV(-40 mV songbird example)
	b	0
	v_{peak}	30 mV
	v_{reset}	-65 mV
	a	0.01 (0.002, songbird example)
	d	200 (100, songbird example)
	I_{Bias}	1000 pA
	k	2.5
	τ_R	2 ms
	τ_D	20 ms
Theta Model	I_{Bias}	0
	τ_R	2
	τ_D	20
LIF Model	τ_m	10 ms
	τ_{ref}	2 ms
	v_{reset}	-65 mV
	v_t	-40 mV
	I_{Bias}	-39 pA

Table 1: The parameters used for the Izhikevich, theta, and LIF neuron models, unless otherwise stated

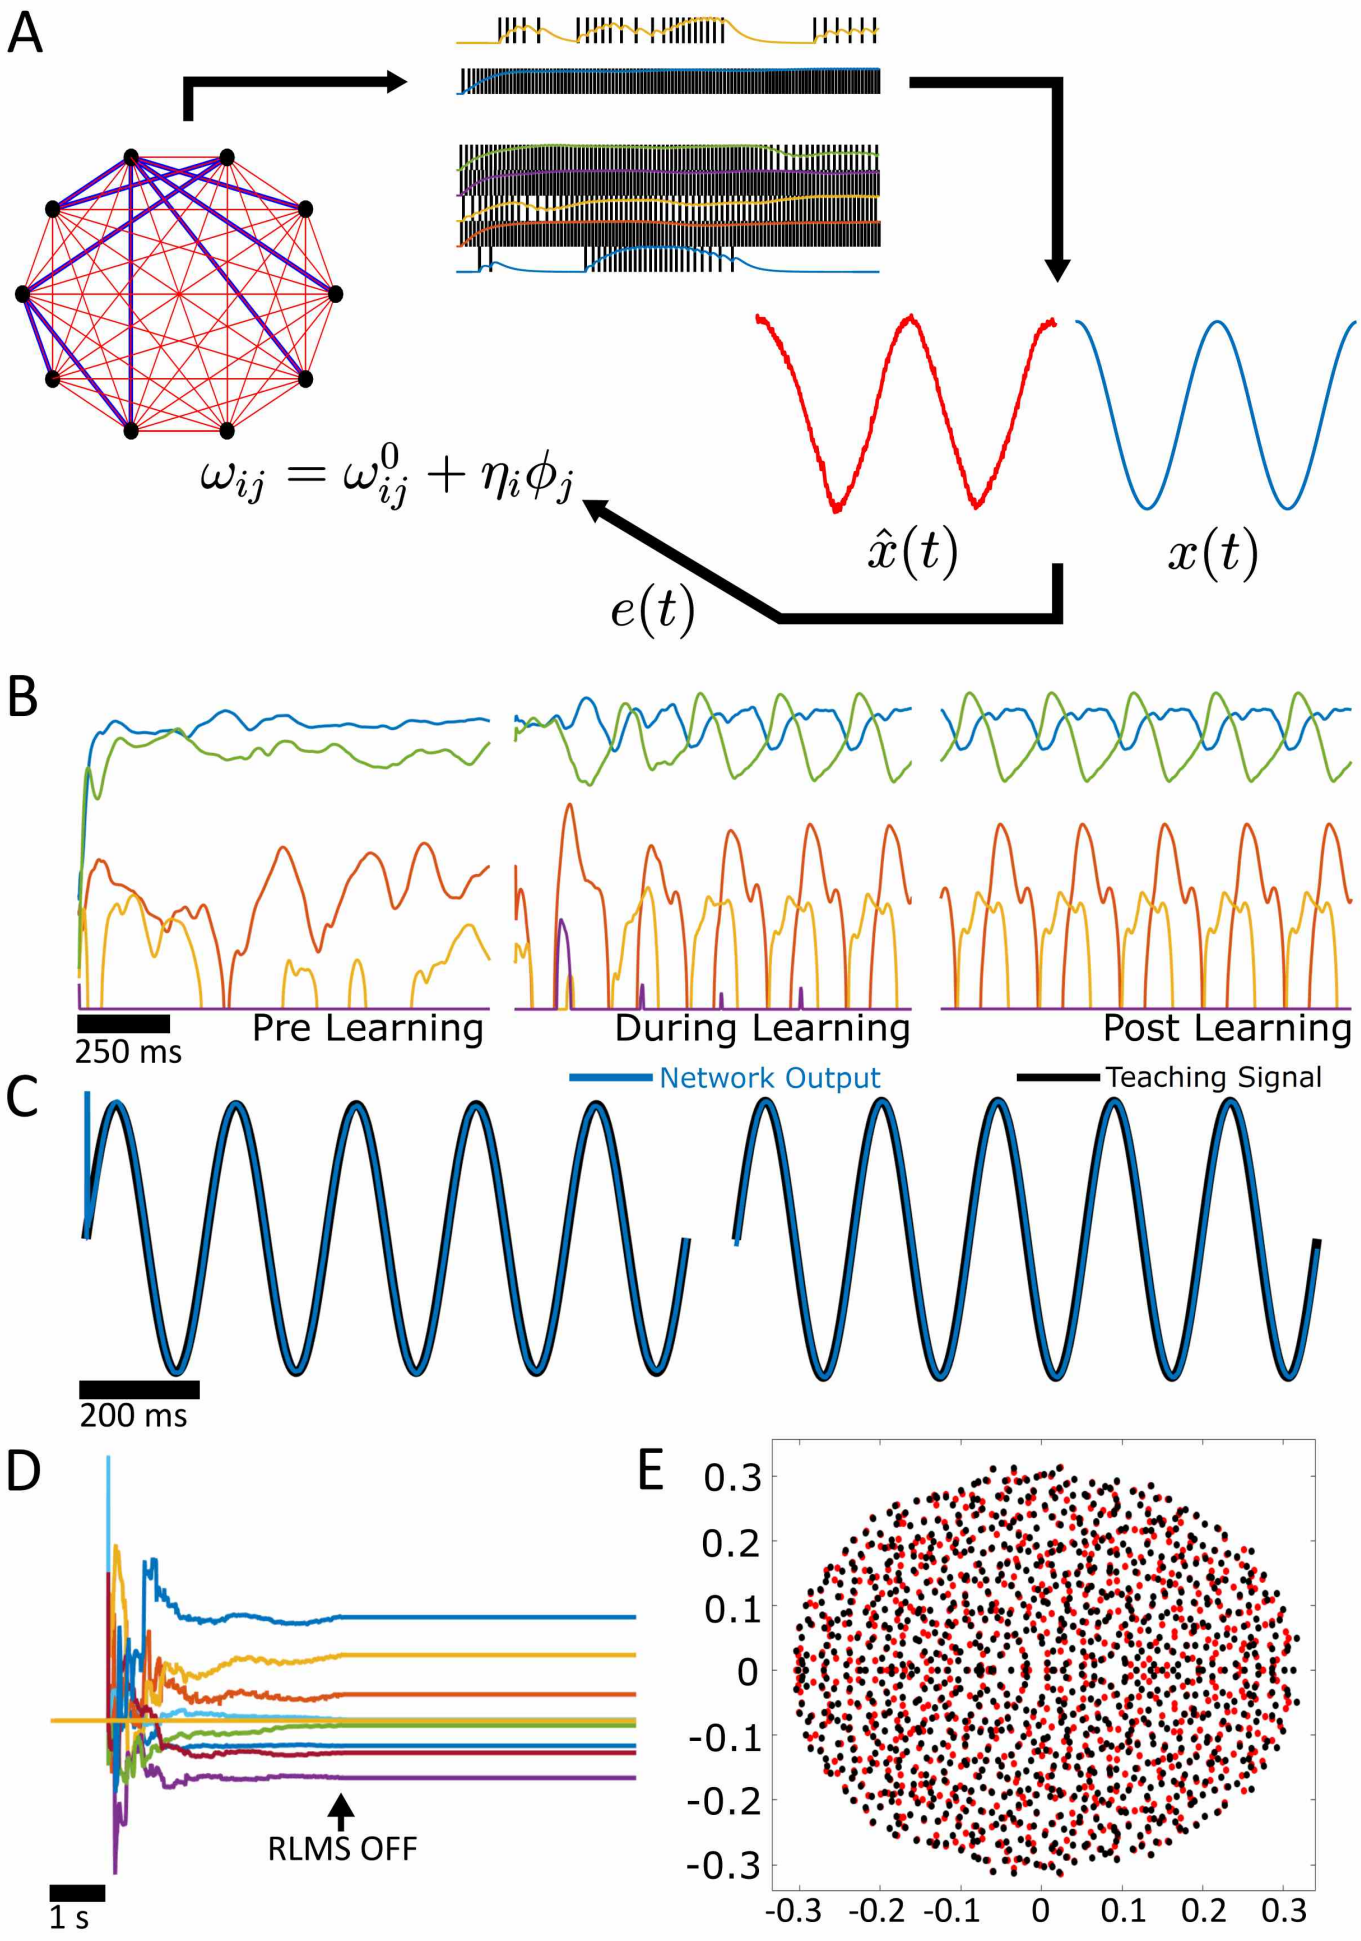
A
$\omega_{ij} = \omega_{ij}^0 + \eta_i \phi_j$
$\hat{x}(t)$
$x(t)$
$e(t)$
B
Pre Learning
During Learning
Post Learning
250 ms
C
Network Output
Teaching Signal
200 ms
D
RLMS OFF
1 s
E
0.3
0.2
0.1
0
-0.1
-0.2
-0.3
-0.3 -0.2 -0.1 0 0.1 0.2 0.3

Figure 1

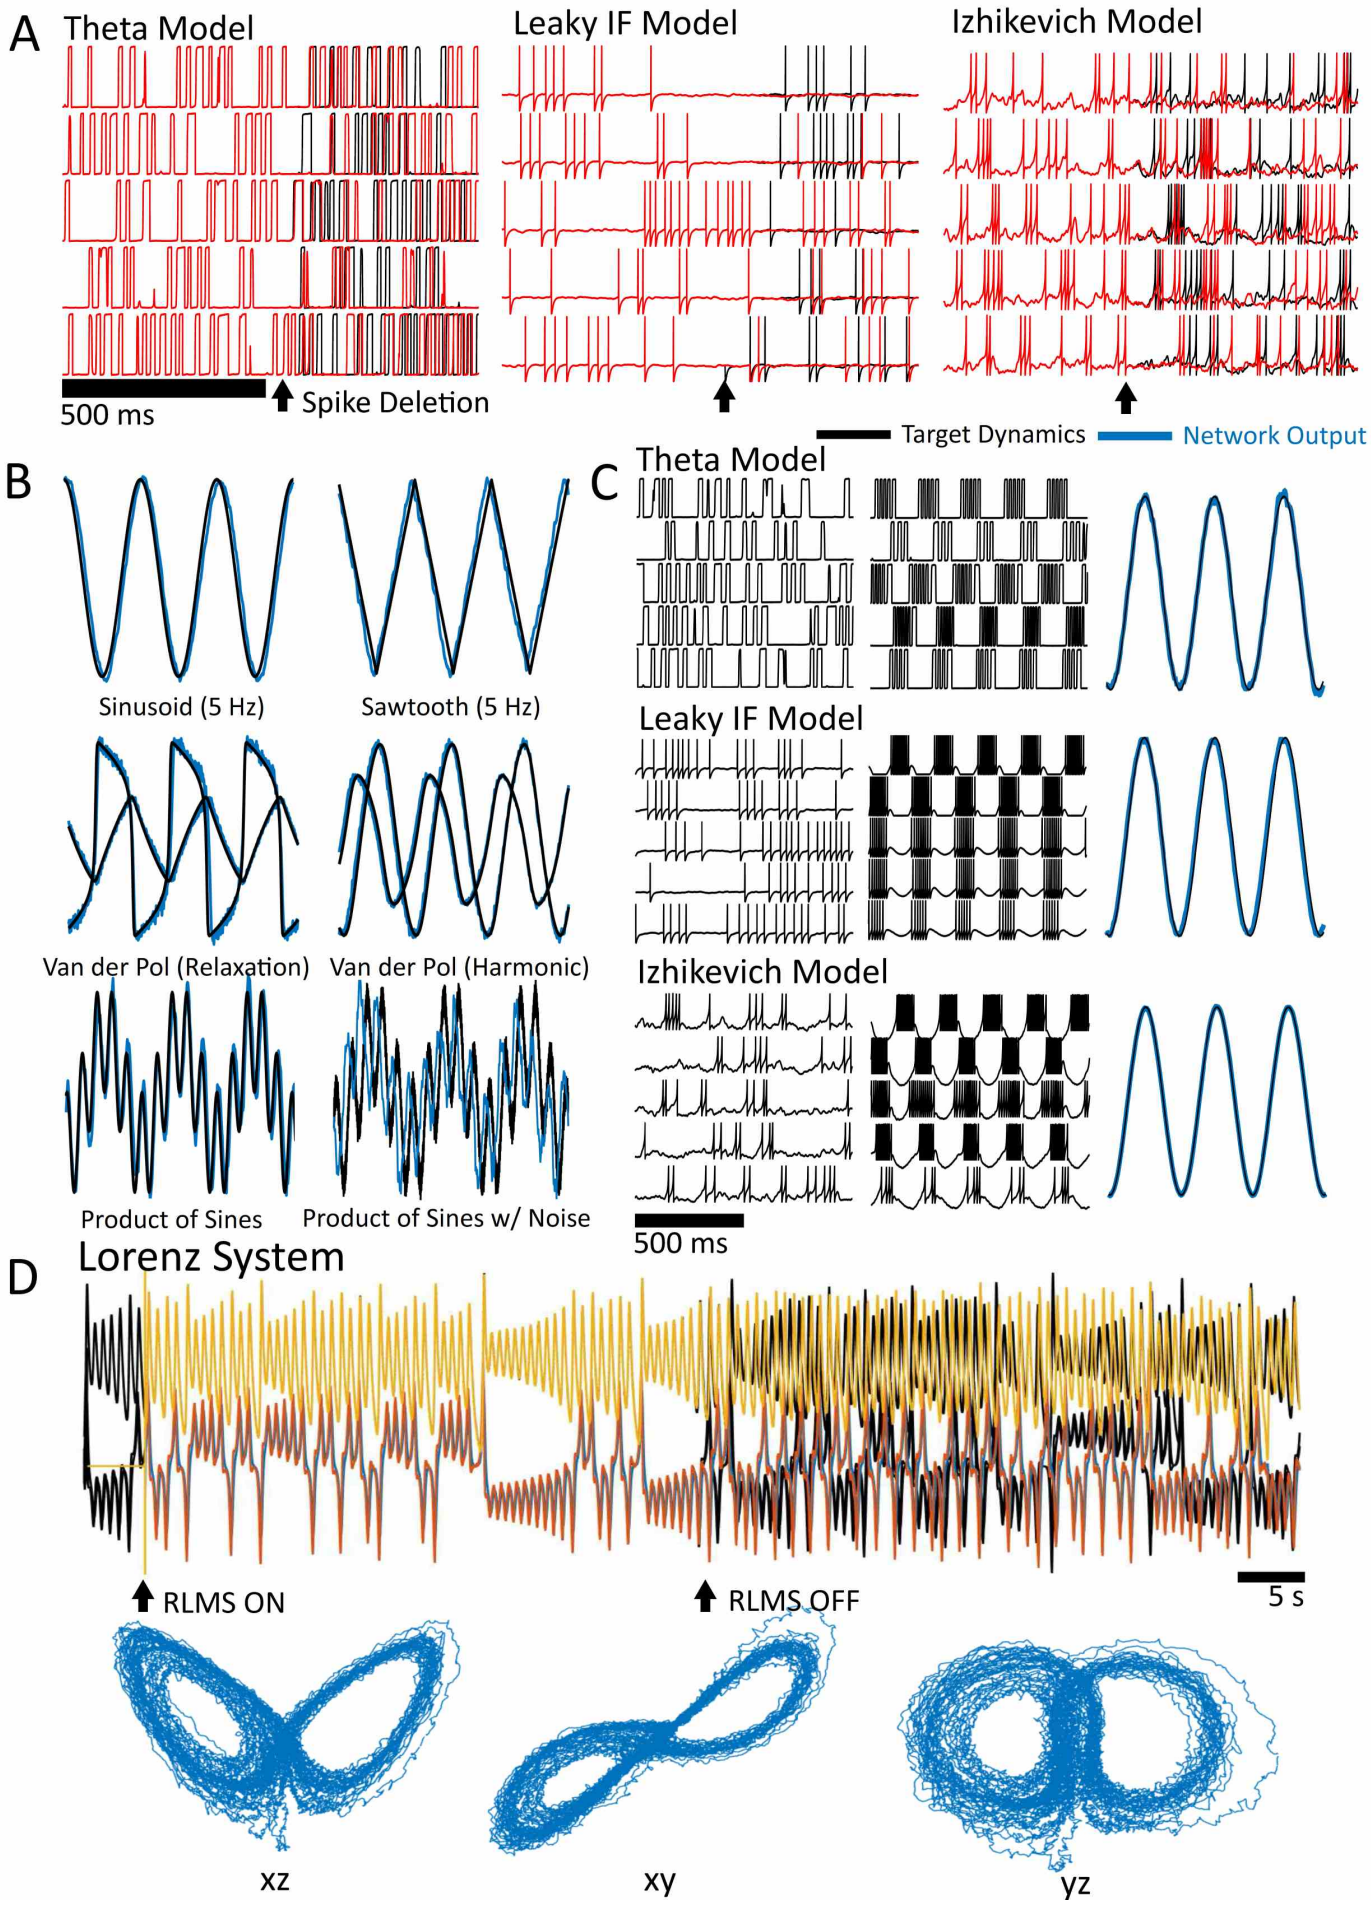
A
Theta Model
Leaky IF Model
Izhikevich Model
500 ms
Spike Deletion
Target Dynamics
Network Output
B
Sinusoid (5 Hz)
Sawtooth (5 Hz)
Van der Pol (Relaxation)
Van der Pol (Harmonic)
Product of Sines
Product of Sines w/ Noise
C
Theta Model
Leaky IF Model
Izhikevich Model
500 ms
D
Lorenz System
RLMS ON
RLMS OFF
5 s
xz
xy
yz

Figure 2

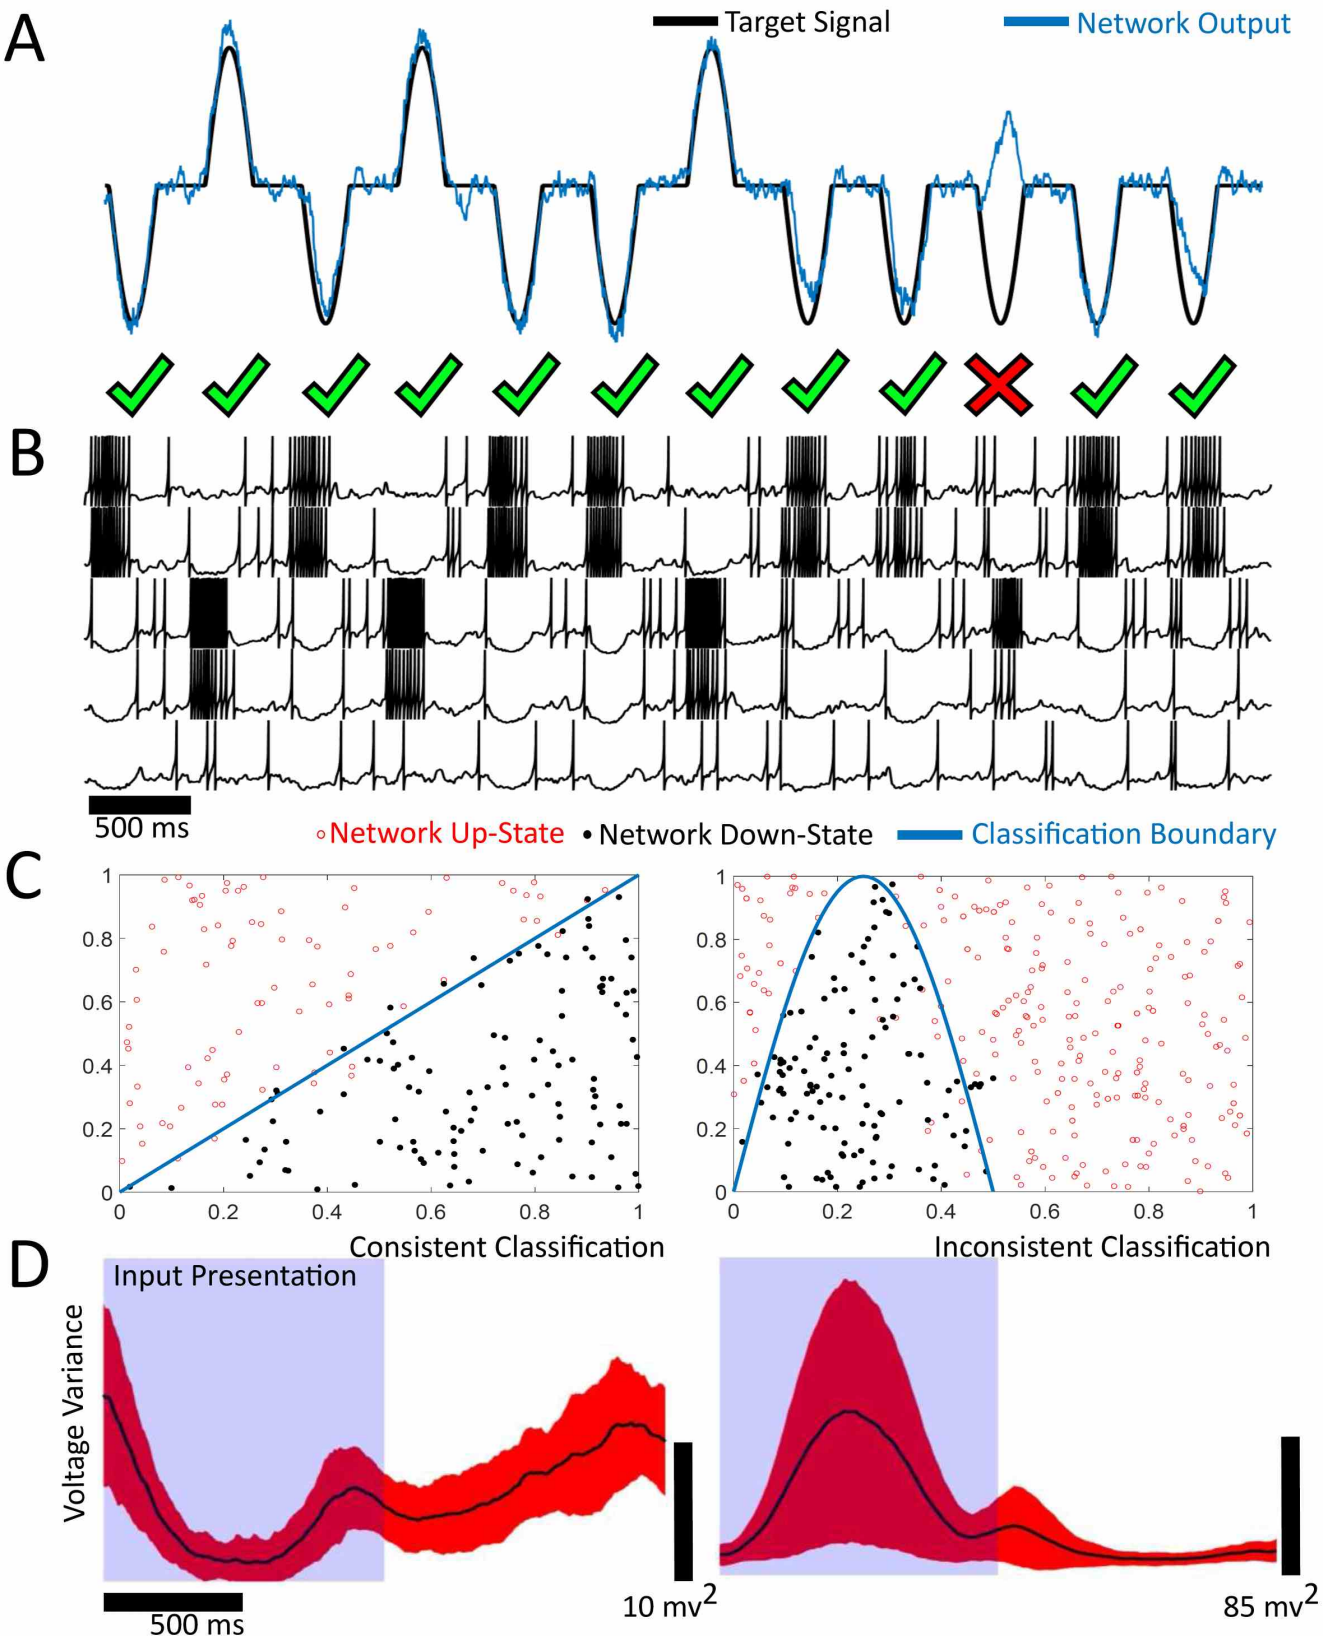
A
Target Signal
Network Output
B
500 ms
Network Up-State
Network Down-State
Classification Boundary
C
Consistent Classification
Inconsistent Classification
D
Input Presentation
Voltage Variance
500 ms
10 mv2
85 mv2

Figure 3

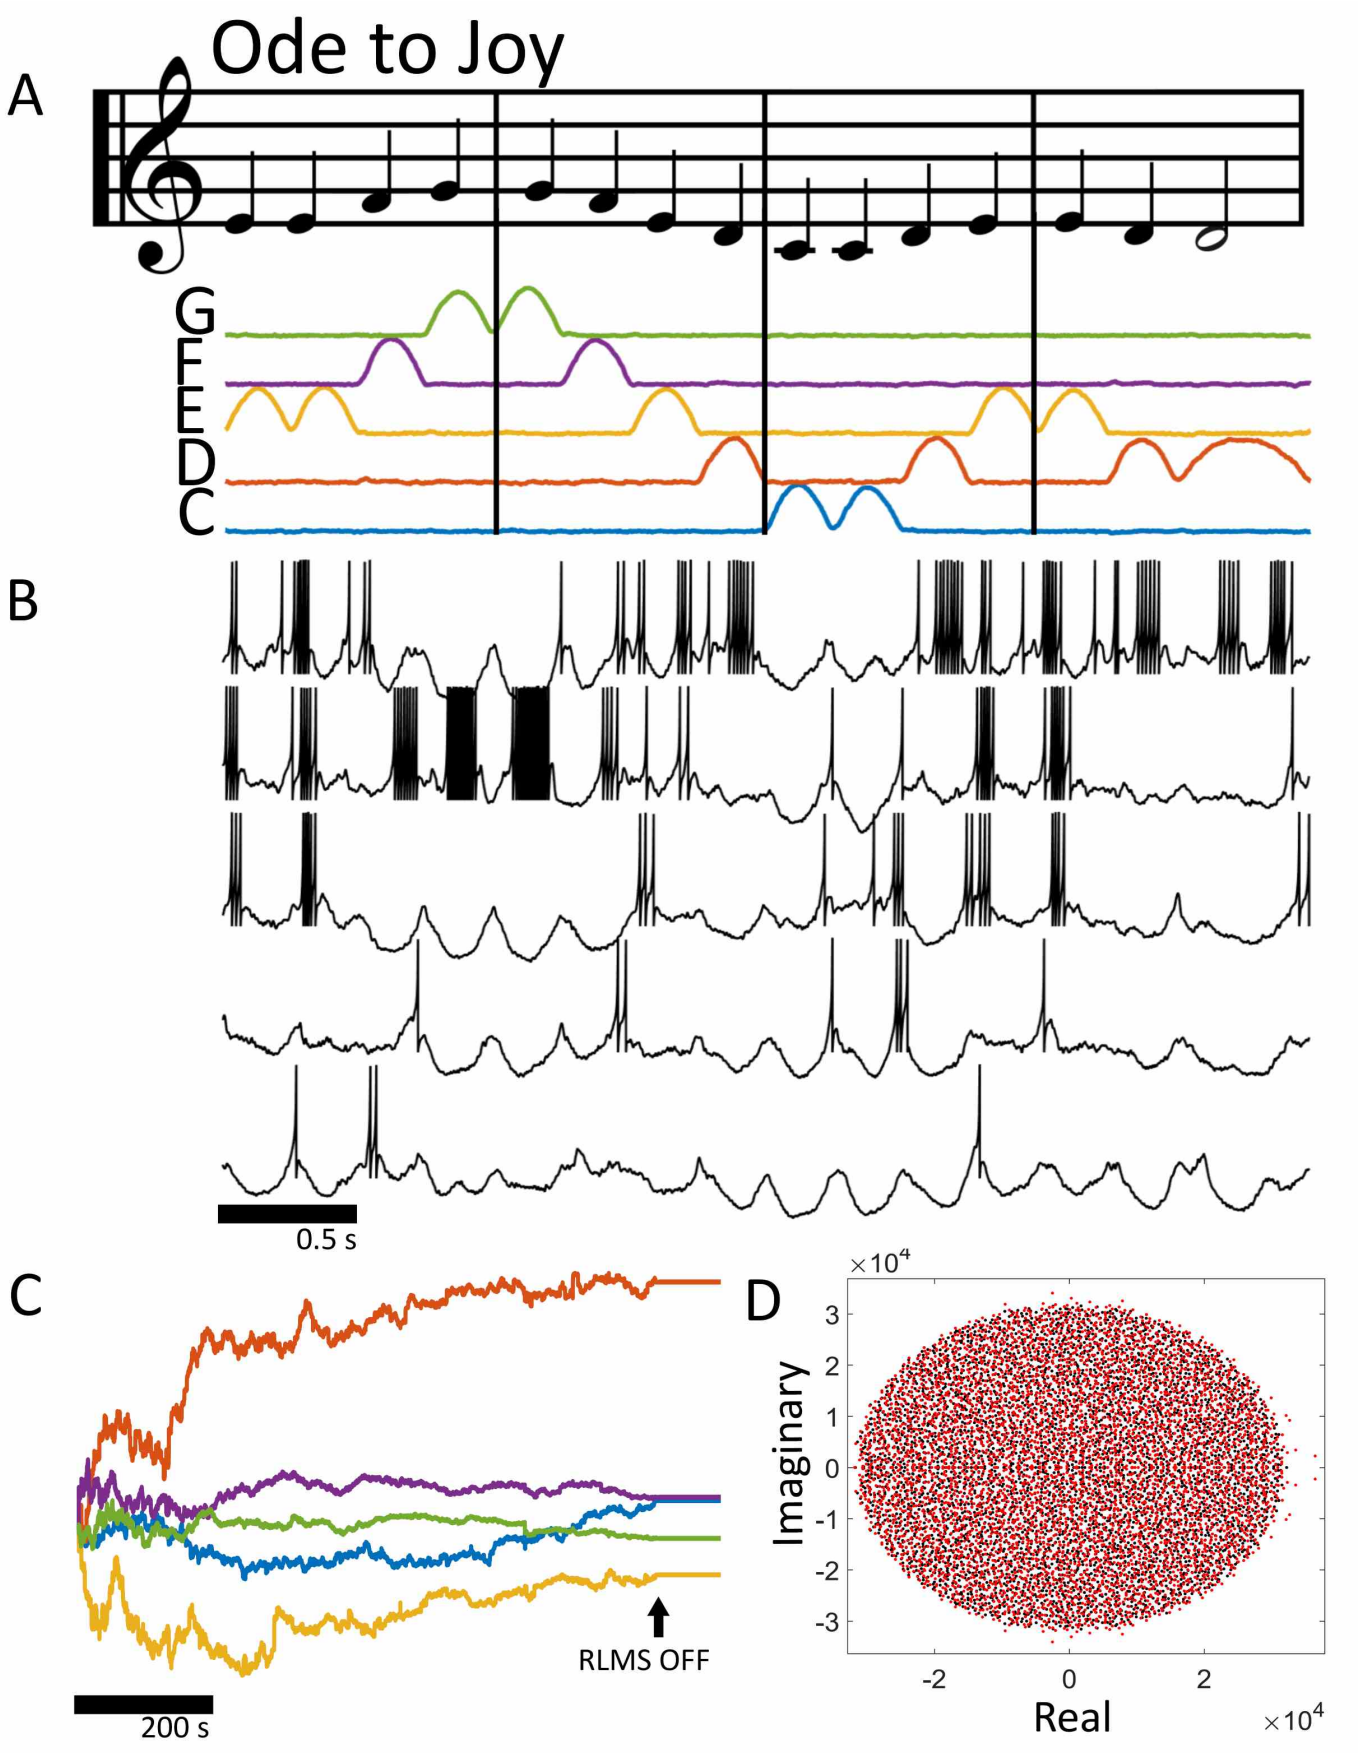
A
Ode to Joy
G
F
E
D
C
B
0.5 s
C
RLMS OFF
200 s
D
×10^4
Imaginary
3
2
1
0
-1
-2
-3
-2
0
2
Real
×10^4

Figure 4

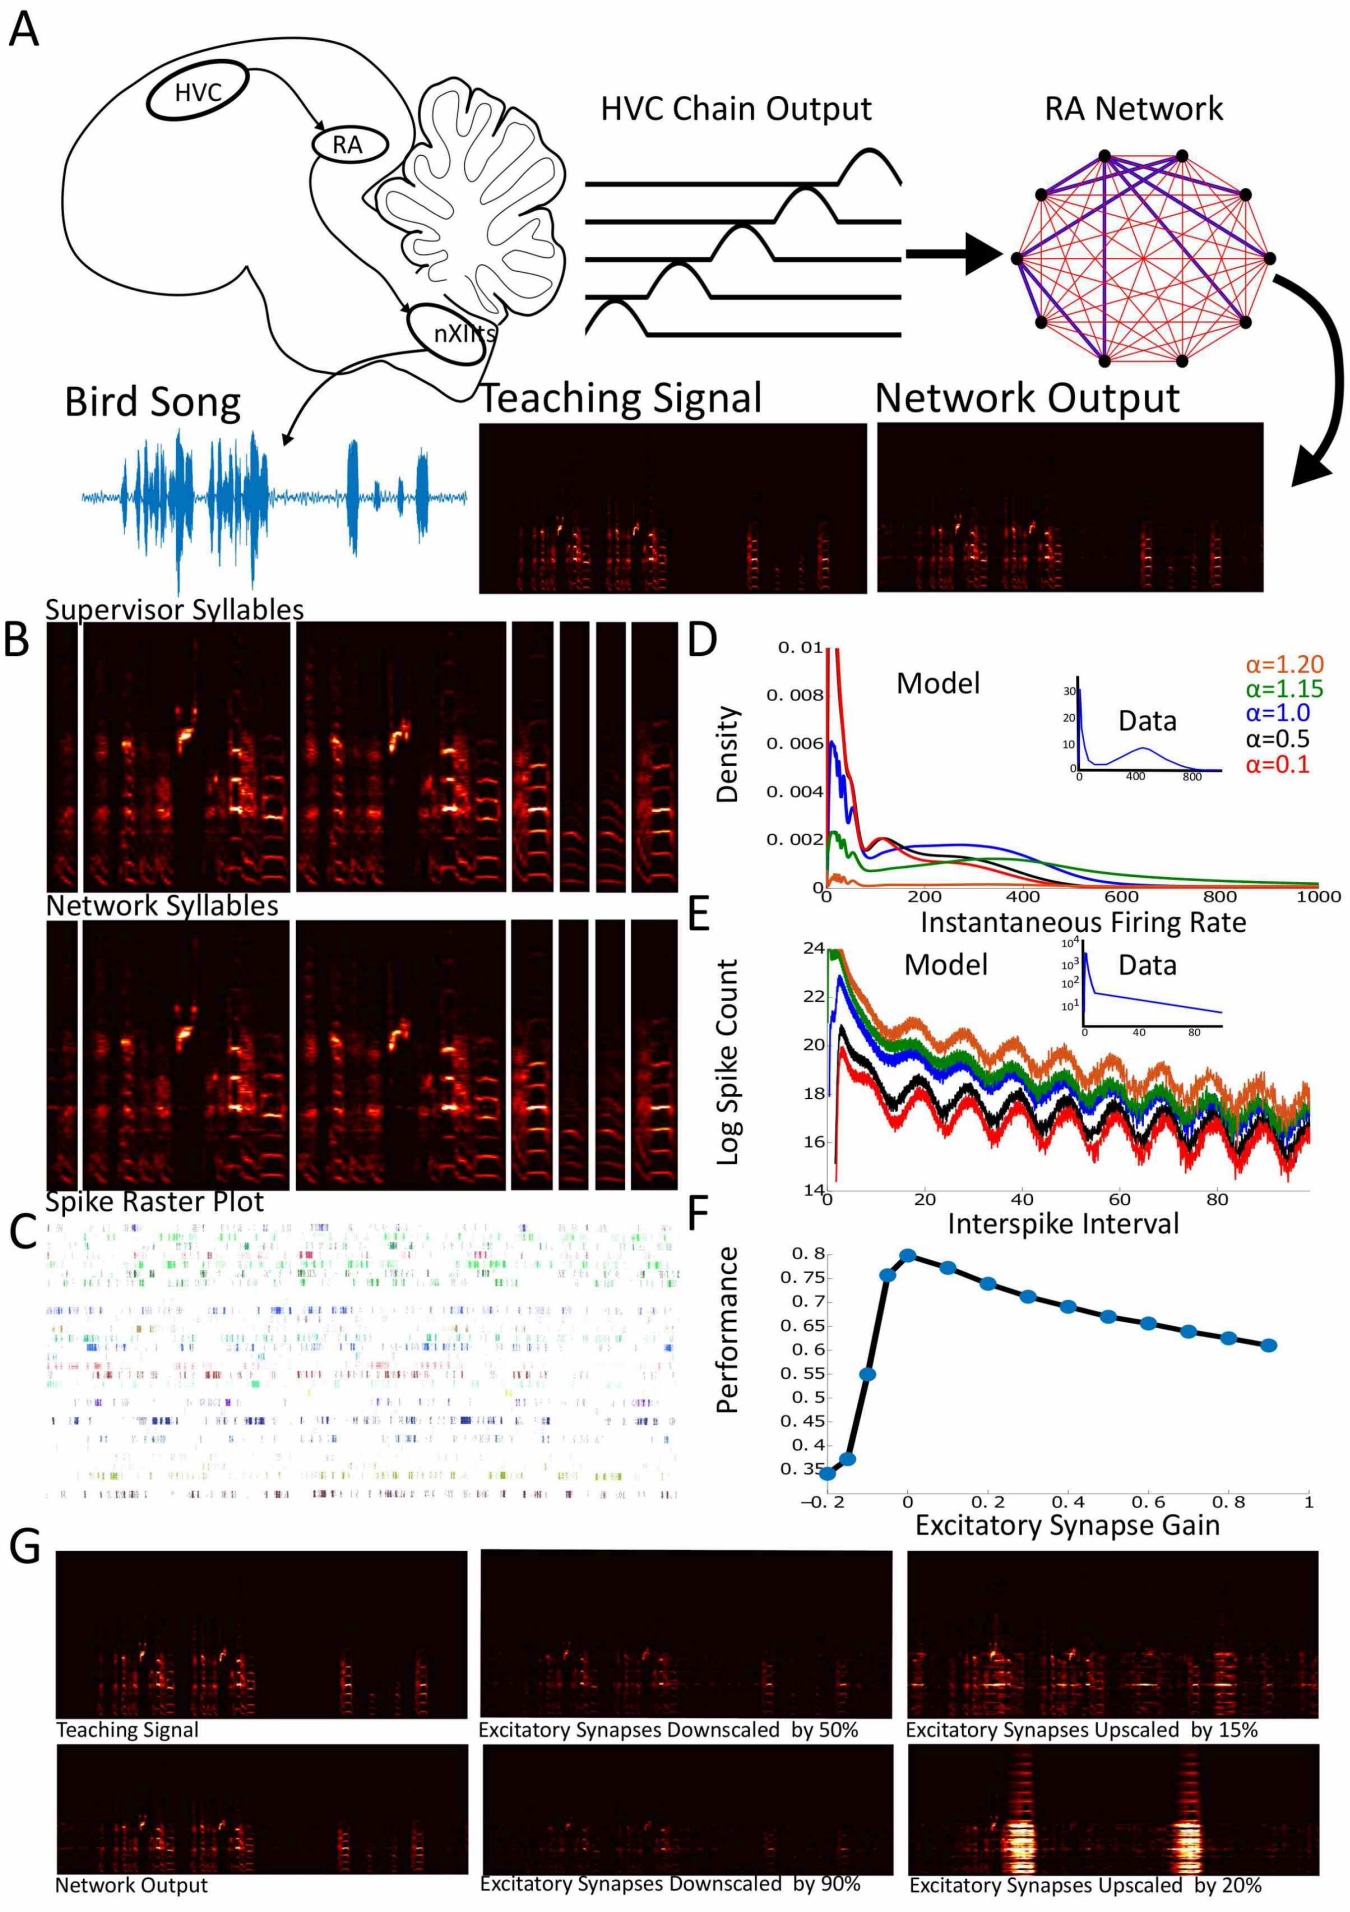
A
HVC
RA
nXIIts
HVC Chain Output
RA Network
Bird Song
Teaching Signal
Network Output
Supervisor Syllables
B
Network Syllables
Spike Raster Plot
C
D
Density
Model
Data
α=1.20
α=1.15
α=1.0
α=0.5
α=0.1
Instantaneous Firing Rate
E
Log Spike Count
Model
Data
F
Interspike Interval
Performance
Excitatory Synapse Gain
G
Teaching Signal
Excitatory Synapses Downscaled by 50%
Excitatory Synapses Upscaled by 15%
Network Output
Excitatory Synapses Downscaled by 90%
Excitatory Synapses Upscaled by 20%

Figure 5

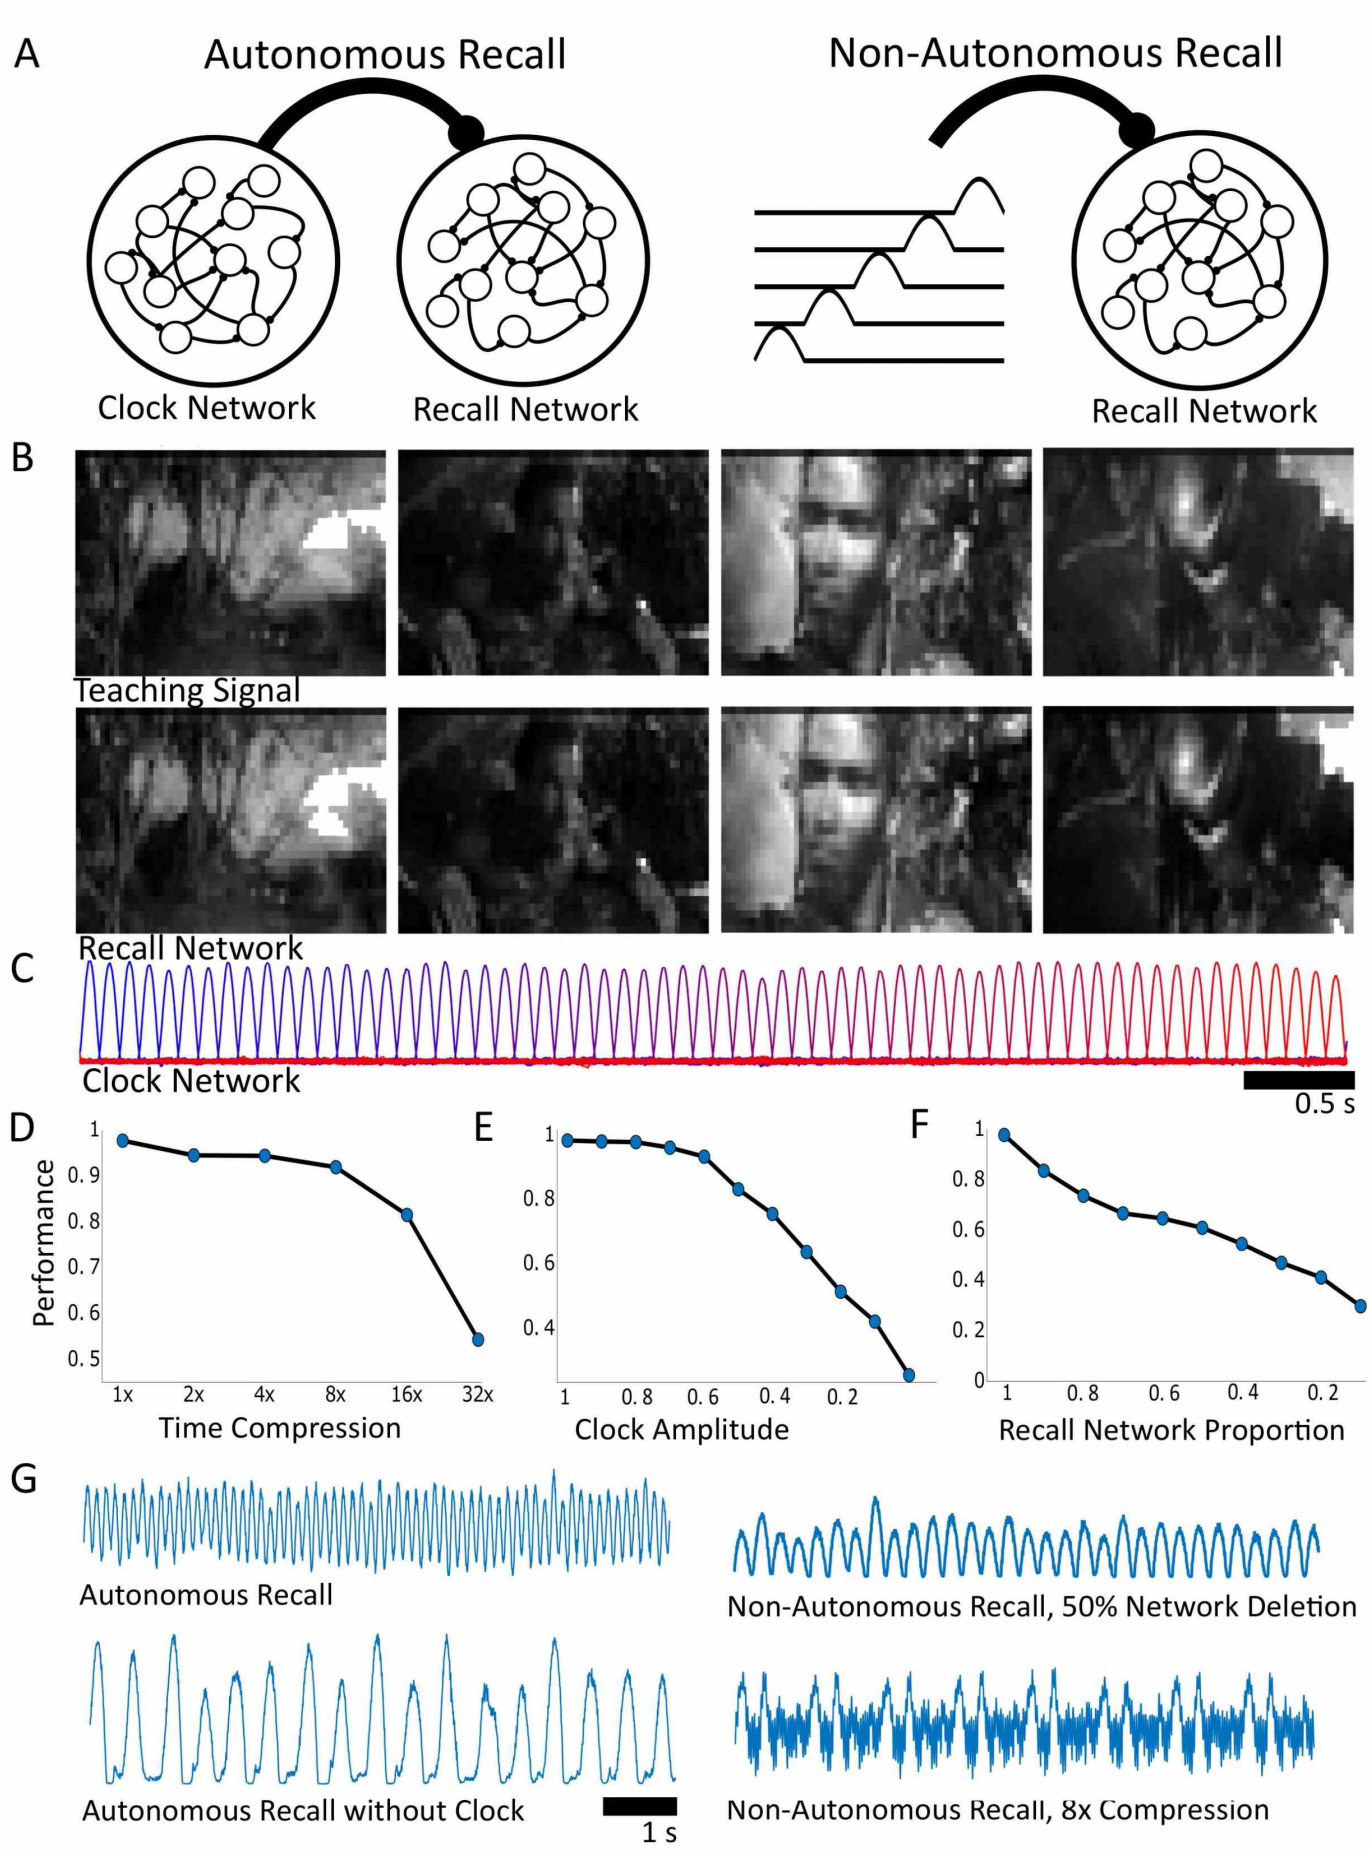
A
Autonomous Recall
Non-Autonomous Recall
Clock Network
Recall Network
Recall Network
B
Teaching Signal
Recall Network
C
Clock Network
0.5 s
D
Performance
1
0.9
0.8
0.7
0.6
0.5
1x
2x
4x
8x
16x
32x
Time Compression
E
1
0.8
0.6
0.4
1
0.8
0.6
0.4
0.2
Clock Amplitude
F
1
0.8
0.6
0.4
0.2
0
1
0.8
0.6
0.4
0.2
Recall Network Proportion
G
Autonomous Recall
Non-Autonomous Recall, 50% Network Deletion
Autonomous Recall without Clock
1 s
Non-Autonomous Recall, 8x Compression

Figure 6

Figure Captions

Figure 1: The FORCE Method with Rate Networks

(A) In the FORCE method, a spiking neural network contains a backbone of static and strong synaptic weights that scale like $1/\sqrt{N}$ to induce network level chaos (blue). A secondary set of weights are added to the weight matrix with the decoders determined through a time optimization procedure. The Recursive Least Mean Squares technique (RLMS) is used in all subsequent simulations. FORCE training requires a supervisor $x(t)$ to estimate with $\hat{x}(t)$. (B) Prior to learning, the firing rates for 5 random neurons from a network of 1000 are in the chaotic regime. The chaos is controlled and converted to steady state oscillations. This allows the network to represent a 5 Hz sinusoidal input (black). After learning, the network (blue) still displays the 5 Hz sinusoidal oscillation as its macroscopic dynamics and the training is successful. The total training time is 5 seconds. (C) The decoders for 20 randomly selected neurons in the network, before ($t < 5$), during ($5 \leq t < 10$) and after ($t \geq 10$) FORCE training. (E) The eigenvalue spectrum for the effective weight matrix before (red) and after (black) FORCE training.

Figure 2: Using Spiking Neural Networks to Mimic Dynamics with FORCE Training

(A) The voltage trace for 5 randomly selected neurons in networks of 2000 integrate-and-fire spiking neurons. The models under consideration are the theta neuron (left), the leaky integrate-and-fire neuron (middle), and the Izhikevich model with spike frequency adaptation (right). For all networks under consideration, a spike was deleted (black arrow) from one of the neurons. This caused the spike train to diverge post-deletion, a clear indication of chaotic behavior. (B) A network of 2000 theta neurons (blue) was initialized in the chaotic regime and trained to mimic different oscillators (black) with FORCE training. The oscillators included the sinusoid, Van der Pol in harmonic and relaxation regimes, a non-smooth sawtooth oscillator, the oscillator formed from taking the product of a pair of sinusoids with 4 Hz and 6 Hz frequencies, and the same product of sinusoids with a Gaussian additive white noise distortion with a standard deviation of 0.05. (C) Three networks of different integrate-and-fire neurons were initialized in the chaotic regime (left), and trained using the FORCE method (center) to mimic a 5 Hz sinusoidal oscillator (right). (D) A network of 5000 theta neurons was trained with the FORCE method to mimic the Lorenz system. RLMS was used to learn the decoders using a 45 second long trajectory of the Lorenz system as a supervisor. RLMS was turned off after 50 seconds with the resulting trajectory and chaotic attractor bearing a strong resemblance to the Lorenz system. The network attractor is shown in three different viewing planes for the 50 seconds post RLMS training.

Figure 3: Using Spiking Neural Networks for Classification with FORCE training

(A) A binary classification task can be performed by a network of 2000 Izhikevich neurons. Up-states correspond to one class, and down-states correspond to another class. The network (blue) is trained with target data (black) using the FORCE method for 800 seconds with an input frequency of 2 Hz on a training set of inputs. The up and down states are formed from the positive and negative components of a 2 Hz sinusoidal function (B) The voltage trace for 3 randomly selected neurons in the network at an identical time to (A). (C) The network correctly classifies data with linear (left) and nonlinear (right) boundaries. The network classifies points as up-states (open red circle) or down states (closed black circle). The x and y axes correspond to the two dimensional input vector. The linear classification task has an accuracy of 0.93 (training time of 800 seconds, 1600 inputs) while the nonlinear task has an accuracy of 0.94 (training time of 800 seconds, 1600 inputs) on a pair of test data sets. (D) The network average voltage variance is (black, standard deviation bars in red) is computed for repeated presentations of a point away from the boundary that can be consistently classified (left) and a point near the boundary that is inconsistently classified for the linear test data (right). The voltage variance only decreases for inputs that the network can consistently classify.

Figure 4: Using Spiking Neural Networks for Pattern Storage and Recall with FORCE Training

(A) The notes in the song Ode to Joy by Beethoven are converted into a continuous 5-component teaching signal. The presence of a note is designated by an up-state in the signal. Quarter notes are the positive portion of a 2 Hz sinusoidal wave form while half notes are represented with the positive part of a 1 Hz sinusoidal wave form. Each component in the teaching signal corresponds to the musical notes C,D,E,F and G. The teaching signal is presented to a network of 5000 Izhikevich neurons continuously from start to finish until the network learns to reproduce the sequence through FORCE training. The network output in (A) is after 900 seconds

of FORCE training, or 225 applications of the target signal. For a 1000 second stretch of simulation time after training, the song was played correctly 205 times comprising 820 seconds of the signal (see supplementary notes for further details). (B) 5 randomly selected neurons in the network and their resulting dynamics after FORCE training. The voltage traces are taken at the same time as the approximant signal in (A). (C) The decoders before ($1 < t < 900$) and after ($t > 900$) FORCE method learning. (D) The resulting eigenvalues for the weight matrix $G\omega_{ij}^0 + Q\boldsymbol{\eta}_i\boldsymbol{\phi}_j$ before (black) and after (red) learning.

Figure 5: Using Spiking Neural Networks for Pattern Storage and Recall: Songbird Singing

(A) A series of up-states constructed from the positive portion of a sinusoidal oscillator with a period of 20 ms are used to model the chain of firing output from RA projection neurons in HVC. These neurons connect to neurons in RA and trigger singing behavior. FORCE training was used to train a network of 1000 Izhikevich neurons to reproduce the spectrogram of a 5 second song clip from an adult zebra finch. (B) The syllables in the song output for the teaching signal (top) and the network (bottom). (C) Spike raster plot for 30 neurons over 5 repetitions of the song. The spike raster plot is aligned with the syllables in (B). (D) The distribution of instantaneous firing rates for the RA network post training. The α parameter corresponds to the up/down regulation factor of the excitatory synapses and corresponds to the different colours in the graphs of (D) and (E). The inset of (D) and (E) is reproduced from [Leonardo and Fee, 2005]. Note that the y-axis of the model and data differ, this is due to normalization of the density functions. (E) The log of the interspike interval histogram. (F) The correlation between the network output and the teaching signal as the excitatory synapses are upscaled (negative x-axis) or downscaled (positive x-axis). (G) The spectrogram for the teaching signal, and the network output under different scaling factors for the excitatory synaptic weights.

Figure 6: Using Spiking Neural Networks for Pattern Storage and Recall: Movie Replay

(A) Two types of networks were trained to recall an 8 second clip from a movie. In autonomous recall networks, both a clock signal and a recall network are simultaneously trained. The clock signal projects onto the recall network similar to how HVC neurons project onto RA neurons in the birdsong circuit. In non-autonomous recall, the network simply receives the clock supervisor without separately training a clock network. (B) The teaching signal is displayed in addition to the network recall at 4 separate times in the movie clip that correspond to distinct scenes. (C) The clock signal consists of 64 up-states generated from the positive component of a sinusoidal oscillator with a period of 250 ms. The colour of the up-states denotes its order in the temporal chain or clock (D) The clock signal is compressed in time which results in speeding up the replay of the movie clip. The time-averaged correlation coefficient between the teaching signal and network output is used to measure performance. This compression was performed on the non-autonomous recall network. (E) The clock amplitude for the network for autonomous recall was reduced. The network was robust to decreasing the amplitude of the clock signal. (F) Neurons in the recall network were removed and recall performance was measured. The recall performance decreases in an approximately linear fashion with the proportion of recall network that is removed. (G) The LFP for the recall networks under a variety of manipulations.

References

[Abbott et al., 2016] Abbott, L., DePasquale, B., and Memmesheimer, R.-M. (2016). Building functional networks of spiking model neurons. *Nature neuroscience*, 19(3):350–355.

[Abbott and van Vreeswijk, 1993] Abbott, L. and van Vreeswijk, C. (1993). Asynchronous states in networks of pulse-coupled oscillators. *Physical Review E*, 48(2):1483.

[Babadi and Abbott, 2016] Babadi, B. and Abbott, L. (2016). Stability and competition in multi-spike models of spike-timing dependent plasticity. *PLoS Comput Biol*, 12(3):e1004750.

[Bi and Poo, 1998] Bi, G.-q. and Poo, M.-m. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of neuroscience*, 18(24):10464–10472.

[Bishop, 2006] Bishop, C. M. (2006). Pattern recognition. *Machine Learning*, 128.

[Boerlin et al., 2013] Boerlin, M., Machens, C. K., and Denève, S. (2013). Predictive coding of dynamical variables in balanced spiking networks. *PLoS Comput Biol*, 9(11):e1003258.

[Bouchard and Brainard, 2016] Bouchard, K. E. and Brainard, M. S. (2016). Auditory-induced neural dynamics in sensory-motor circuitry predict learned temporal and sequential statistics of birdsong. *Proceedings of the National Academy of Sciences*, 113(34):9641–9646.

[Bourdoukan and Denève, 2015] Bourdoukan, R. and Denève, S. (2015). Enforcing balance allows local supervised learning in spiking recurrent networks. In *Advances in Neural Information Processing Systems*, pages 982–990.

[Boyce et al., 2016] Boyce, R., Glasgow, S. D., Williams, S., and Adamantidis, A. (2016). Causal evidence for the role of rem sleep theta rhythm in contextual memory consolidation. *Science*, 352(6287):812–816.

[Bressloff, 2011] Bressloff, P. C. (2011). Spatiotemporal dynamics of continuum neural fields. *Journal of Physics A: Mathematical and Theoretical*, 45(3):033001.

[Buonomano and Merzenich, 1995] Buonomano, D. V. and Merzenich, M. M. (1995). Temporal information transformed into a spatial code by a neural network with realistic properties. *Science*, 267(5200):1028.

[Buzsáki, 2002] Buzsáki, G. (2002). Theta oscillations in the hippocampus. *Neuron*, 33(3):325–340.

[Buzsáki, 2005] Buzsáki, G. (2005). Theta rhythm of navigation: link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus*, 15(7):827–840.

[Buzsáki, 2015] Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, 25(10):1073–1188.

[Churchland et al., 2010] Churchland, M. M., Byron, M. Y., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., Newsome, W. T., Clark, A. M., Hosseini, P., Scott, B. B., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature neuroscience*, 13(3):369–378.

[Churchland et al., 2012] Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., and Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*, 487(7405):51–56.

[Clopath et al., 2010] Clopath, C., Büsing, L., Vasilaki, E., and Gerstner, W. (2010). Connectivity reflects coding: a model of voltage-based stdp with homeostasis. *Nature neuroscience*, 13(3):344–352.

[Coombes, 2005] Coombes, S. (2005). Waves, bumps, and patterns in neural field theories. *Biological cybernetics*, 93(2):91–108.

[Crandall and Nick, 2014] Crandall, S. R. and Nick, T. A. (2014). The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *CRCNS.org*.

[DePasquale et al., 2016] DePasquale, B., Churchland, M. M., and Abbott, L. (2016). Using firing-rate dynamics to train recurrent networks of spiking model neurons. *arXiv preprint arXiv:1601.07620*.

[Diba and Buzsáki, 2007] Diba, K. and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature neuroscience*, 10(10):1241–1242.

[Dominey, 1995] Dominey, P. F. (1995). Complex sensory-motor sequence learning based on recurrent state representation and reinforcement learning. *Biological cybernetics*, 73(3):265–274.

[Eliasmith and Anderson, 2004] Eliasmith, C. and Anderson, C. H. (2004). *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. MIT press.

[Eliasmith et al., 2012] Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., and Rasmussen, D. (2012). A large-scale model of the functioning brain. *science*, 338(6111):1202–1205.

[Euston et al., 2007] Euston, D. R., Tatsuno, M., and McNaughton, B. L. (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *science*, 318(5853):1147–1150.

[Finn et al., 2007] Finn, I. M., Priebe, N. J., and Ferster, D. (2007). The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *Neuron*, 54(1):137–152.

[Givens and Olton, 1995] Givens, B. and Olton, D. S. (1995). Bidirectional modulation of scopolamine-induced working memory impairments by muscarinic activation of the medial septal area. *Neurobiology of learning and memory*, 63(3):269–276.

[Graupner and Brunel, 2012] Graupner, M. and Brunel, N. (2012). Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. *Proceedings of the National Academy of Sciences*, 109(10):3991–3996.

[Grosmark and Buzsáki, 2016] Grosmark, A. D. and Buzsáki, G. (2016). Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science*, 351(6280):1440–1443.

[Hahnloser et al., 2002] Hahnloser, R. H., Kozhevnikov, A. A., and Fee, M. S. (2002). An ultra-sparse code underliesthe generation of neural sequences in a songbird. *Nature*, 419(6902):65–70.

[Harish and Hansel, 2015] Harish, O. and Hansel, D. (2015). Asynchronous rate chaos in spiking neuronal circuits. *PLoS Comput Biol*, 11(7):e1004266.

[Haykin, 2009] Haykin, S. S. (2009). *Neural networks and learning machines*, volume 3. Pearson Upper Saddle River, NJ, USA:.

[Jaeger, 2001] Jaeger, H. (2001). The "echo state" approach to analysing and training recurrent neural networks-with an erratum note. *Bonn, Germany: German National Research Center for Information Technology GMD Technical Report*, 148:34.

[Jarvis et al., 2000] Jarvis, E. D., Ribeiro, S., Da Silva, M. L., Ventura, D., Vielliard, J., and Mello, C. V. (2000). Behaviourally driven gene expression reveals song nuclei in hummingbird brain. *Nature*, 406(6796):628–632.

[Klimesch, 1999] Klimesch, W. (1999). Eeg alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. *Brain research reviews*, 29(2):169–195.

[Knight, 1972] Knight, B. W. (1972). Dynamics of encoding in a population of neurons. *The Journal of general physiology*, 59(6):734–766.

[Konishi, 1985] Konishi, M. (1985). Birdsong: from behavior to neuron. *Annual review of neuroscience*, 8(1):125–170.

[Kozhevnikov and Fee, 2007] Kozhevnikov, A. A. and Fee, M. S. (2007). Singing-related activity of identified hvc neurons in the zebra finch. *Journal of neurophysiology*, 97(6):4271–4283.

[Laje and Buonomano, 2013] Laje, R. and Buonomano, D. V. (2013). Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nature neuroscience*, 16(7):925–933.

[Lengyel et al., 2005] Lengyel, M., Huhn, Z., and Érdi, P. (2005). Computational theories on the function of theta oscillations. *Biological cybernetics*, 92(6):393–408.

[Leonardo and Fee, 2005] Leonardo, A. and Fee, M. S. (2005). Ensemble coding of vocal control in birdsong. *The Journal of Neuroscience*, 25(3):652–661.

[Li et al., 2016] Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. *Nature*, 532(7600):459–464.

[Lisman and Jensen, 2013] Lisman, J. E. and Jensen, O. (2013). The theta-gamma neural code. *Neuron*, 77(6):1002–1016.

[London et al., 2010] London, M., Roth, A., Beeren, L., Häusser, M., and Latham, P. E. (2010). Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature*, 466(7302):123–127.

[Long et al., 2010] Long, M. A., Jin, D. Z., and Fee, M. S. (2010). Support for a synaptic chain model of neuronal sequence generation. *Nature*, 468(7322):394–399.

[Lourenço, 2006] Lourenço, C. (2006). Dynamical reservoir properties as network effects. In *ESANN*, pages 503–508.

[Lubenov and Siapas, 2009] Lubenov, E. V. and Siapas, A. G. (2009). Hippocampal theta oscillations are travelling waves. *Nature*, 459(7246):534–539.

[LukošEvičIus and Jaeger, 2009] LukošEvičIus, M. and Jaeger, H. (2009). Reservoir computing approaches to recurrent neural network training. *Computer Science Review*, 3(3):127–149.

[Lukoševičius et al., 2012] Lukoševičius, M., Jaeger, H., and Schrauwen, B. (2012). Reservoir computing trends. *KI-Künstliche Intelligenz*, 26(4):365–371.

[Maass, 2010] Maass, W. (2010). Liquid state machines: motivation, theory, and applications. *Computability in context: computation and logic in the real world*, pages 275–296.

[Maass et al., 2002] Maass, W., Natschläger, T., and Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural computation*, 14(11):2531–2560.

[Maass et al., 2004] Maass, W., Natschläger, T., and Markram, H. (2004). Fading memory and kernel properties of generic cortical microcircuit models. *Journal of Physiology-Paris*, 98(4):315–330.

[Mello et al., 2015] Mello, G. B., Soares, S., and Paton, J. J. (2015). A scalable population code for time in the striatum. *Current Biology*, 25(9):1113–1122.

[Mooney, 2000] Mooney, R. (2000). Different subthreshold mechanisms underlie song selectivity in identified hvc neurons of the zebra finch. *The Journal of Neuroscience*, 20(14):5420–5436.

[Motanis and Buonomano, 2015] Motanis, H. and Buonomano, D. V. (2015). Neural coding: Time contraction and dilation in the striatum. *Current Biology*, 25(9):R374–R376.

[Nádasdy et al., 1999] Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., and Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *The Journal of Neuroscience*, 19(21):9497–9507.

[Nesse et al., 2008] Nesse, W. H., Borisyuk, A., and Bressloff, P. C. (2008). Fluctuation-driven rhythmogenesis in an excitatory neuronal network with slow adaptation. *Journal of computational neuroscience*, 25(2):317–333.

[Nottebohm et al., 1976] Nottebohm, F., Stokes, T. M., and Leonard, C. M. (1976). Central control of song in the canary, serinus canarius. *Journal of Comparative Neurology*, 165(4):457–486.

[Nykamp and Tranchina, 2001] Nykamp, D. Q. and Tranchina, D. (2001). A population density approach that facilitates large-scale modeling of neural networks: extension to slow inhibitory synapses. *Neural computation*, 13(3):511–546.

[Omurtag et al., 2000a] Omurtag, A., Knight, B., and Sirovich, L. (2000a). Dynamics of neuronal populations: The equilibrium solution. *SIAM Journal on Applied Mathematics*, 60(6):2009–2028.

[Omurtag et al., 2000b] Omurtag, A., Knight, B. W., and Sirovich, L. (2000b). On the simulation of large populations of neurons. *Journal of computational neuroscience*, 8(1):51–63.

[Ostojic, 2014] Ostojic, S. (2014). Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons. *Nature neuroscience*, 17(4):594–600.

[Ozturk and Principe, 2005] Ozturk, M. C. and Principe, J. C. (2005). Computing with transiently stable states. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 3, pages 1467–1472. IEEE.

[Pfister and Gerstner, 2006] Pfister, J.-P. and Gerstner, W. (2006). Triplets of spikes in a model of spike timing-dependent plasticity. *The Journal of neuroscience*, 26(38):9673–9682.

[Rajan and Abbott, 2006] Rajan, K. and Abbott, L. (2006). Eigenvalue spectra of random matrices for neural networks. *Physical review letters*, 97(18):188104.

[Rajan et al., 2016] Rajan, K., Harvey, C. D., and Tank, D. W. (2016). Recurrent network models of sequence generation and memory. *Neuron*, 90(1):128–142.

[Schliebs et al., 2011] Schliebs, S., Mohemmed, A., and Kasabov, N. (2011). Are probabilistic spiking neural networks suitable for reservoir computing? In *Neural Networks (IJCNN), The 2011 International Joint Conference on*, pages 3156–3163. IEEE.

[Schrauwen et al., 2007] Schrauwen, B., Verstraeten, D., and Van Campenhout, J. (2007). An overview of reservoir computing: theory, applications and implementations. In *Proceedings of the 15th European Symposium on Artificial Neural Networks. p. 471-482 2007*, pages 471–482.

[Schwemmer et al., 2015] Schwemmer, M. A., Fairhall, A. L., Denéve, S., and Shea-Brown, E. T. (2015). Constructing precisely computing networks with biophysical spiking neurons. *The Journal of Neuroscience*, 35(28):10112–10134.

[Sirovich et al., 2006] Sirovich, L., Omurtag, A., and Lubliner, K. (2006). Dynamics of neural populations: Stability and synchrony. *Network: Computation in neural systems*, 17(1):3–29.

[Strogatz, 2000] Strogatz, S. H. (2000). From kuramoto to crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena*, 143(1):1–20.

[Strogatz and Mirollo, 1991] Strogatz, S. H. and Mirollo, R. E. (1991). Stability of incoherence in a population of coupled oscillators. *Journal of Statistical Physics*, 63(3-4):613–635.

[Sussillo and Abbott, 2009] Sussillo, D. and Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63(4):544–557.

[Sussillo et al., 2015] Sussillo, D., Churchland, M. M., Kaufman, M. T., and Shenoy, K. V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nature neuroscience*, 18(7):1025–1033.

[Thalmeier et al., 2015] Thalmeier, D., Uhlmann, M., Kappen, H. J., and Memmesheimer, R.-M. (2015). Learning universal computations with spikes. *arXiv preprint arXiv:1505.07866*.

[van Vreeswijk, 1996] van Vreeswijk, C. (1996). Partial synchronization in populations of pulse-coupled oscillators. *Physical Review E*, 54(5):5522.

[Van Vreeswijk et al., 1994] Van Vreeswijk, C., Abbott, L., and Ermentrout, G. B. (1994). When inhibition not excitation synchronizes neural firing. *Journal of computational neuroscience*, 1(4):313–321.

[van Vreeswijk and Sompolinsky, 1996] van Vreeswijk, C. and Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724.

[Vogels et al., 2011] Vogels, T., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science*, 334(6062):1569–1573.

[Wilson and Cowan, 1972] Wilson, H. R. and Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical journal*, 12(1):1.

[Wojcik and Kaminski, 2004] Wojcik, G. M. and Kaminski, W. A. (2004). Liquid state machine built of hodgkin–huxley neurons and pattern recognition. *Neurocomputing*, 58:245–251.

Supplementary Material

S1: Comparison with The Neural Engineering Framework

Like the FORCE method, the Neural Engineering Framework (NEF) is a technique for creating functional, spiking networks of neurons. These two approaches can yield weight matrices with similar structure depending on the parameter regime. Thus, we will demonstrate how NEF generated networks are derived and the relationship between the weight matrices generated by these two different approaches.

The NEF approach generates weight matrices to approximate the dynamical system $\dot{\boldsymbol{x}} = G(\boldsymbol{x})$, a system of autonomous differential equations. This is achieved through assuming that the weight matrix is a low rank weight matrix of the form:

$$\omega_{ij} = \boldsymbol{\eta}_i \boldsymbol{\phi}_j^T \quad \text{(NEF Weight Matrix)} \tag{19}$$

where $\boldsymbol{\eta}_i$ is a quantity referred to as the neurons encoder and sets the preferred orientation vector for neuron i. The quantity $\boldsymbol{\phi}_j$ is referred to as the neuronal decoder and is used to set the dynamics of the network through a static optimization problem. Note that if we consider FORCE trained matrices, we have the following:

$$\omega_{ij} = G\omega_{ij}^0 + Q\boldsymbol{\eta}_i \boldsymbol{\phi}_j^T \quad \text{(FORCE Weight Matrix)} \tag{20}$$

Thus, we arrive at NEF assumption for weight matrices in the limit that $Q \gg G$, and redefinition of $\boldsymbol{\eta}_i$. The weight matrix for NEF networks is assumed to have a rank that is equal to the dimensionality of the target dynamical system. This is in contrast to FORCE generated weight matrices which are always full rank. The decoders can be solved for with a static least squares or L_2 optimization over the phase space of x when one approximates the neuronal spiking with rates. These networks can be improved upon by incorporating spike timing and neuronal correlations by applying a second L_2 optimization over time after the network has been simulated. The first static phase space optimization yields the network that approximates the dynamical system $\dot{\boldsymbol{x}} = G(\boldsymbol{x})$ while the second optimization extracts out the best approximant from the spiking network.

The procedure for generating spiking networks using the NEF method is shown in Figure S1 for a network of 5000 theta-neurons with $p = 0.15$ sparse coupling between the connections where the intended dynamics is a Lorenz system. The default NEF approximant is shown in Figure S1B, in addition to the trajectory that uses neuronal correlations in the spiking network optimally. Due to the chaotic nature of the signal, the error signal in the second optimization procedure is computed by taking the trajectory generated by NEF network, and smoothing out the high frequency components from spiking with a moving average. As more of the signal is used in the optimization procedure, the resulting normalized L_2 error decreases substantially for the remainder of the signal ($t > T$). The eigenvalue spectrum of the resulting weight matrix is shown in Figure S1D, which shows a cloud of eigenvalues bounded by a circle, in addition to three dominant eigenvalues that correspond to the three-dimensional dynamics of the Lorenz system. The NEF approach can generate networks that mimic other dynamical systems, such as the Van der Pol system, shown in Figure S1E, in different parameter regimes and the normal form for the cusp bifurcation, $\dot{x} = \mu x - x^3$.

NEF Derivation

Consider a network of integrate-and-fire neurons receiving currents of the form $I_i = \sum_{j=1}^N \omega_{ij} r_j(t) + \beta_i$ where β_i is the heterogeneously distributed bias current, and r_j is the postsynaptically filtered spike train for the jth neuron, with each neuron having some known frequency response curve, $f(I)$. Like the FORCE method, the NEF uses a linear combination of synaptically filtered spike trains as the approximant:

$$\hat{x}(t) = \sum_{j=1}^{N} \boldsymbol{\phi}_j r_j(t) \tag{21}$$

where $r_j(t)$ is the filtered spike train for the jth neuron. For example, the single exponential filter is given by

$$r_j' = -\frac{r_j}{\tau_s} + \sum_{t_{jk}<t} \delta(t - t_{jk}) \tag{22}$$

where t_{jk} is the kth spike fired by the jth neuron. The ϕ_j are determined by an optimization problem. This problem is derived by first determining the dynamics of the approximant in terms of spikes, and subsequently

replacing the spike train with rates. For example, for the single exponential synapse we have the following:

$$\begin{aligned}
\frac{d\hat{x}}{dt} &= \sum_{j=1}^{N} \boldsymbol{\phi}_j \left(-\frac{r_j}{\tau_s} + \frac{1}{\tau_s} \sum_{t_{jk}<t} \delta(t - t_{jk}) \right) && (23)\\
&= -\frac{\hat{x}}{\tau_s} + \sum_{j=1}^{N} \sum_{t_{jk}<t} \boldsymbol{\phi}_j \delta(t - t_{jk}) && (24)\\
&\approx -\frac{\hat{x}}{\tau_s} + \frac{1}{\tau_s} \sum_{j=1}^{N} \phi_j f(\boldsymbol{\eta}_j \hat{x} + \beta_j) = \hat{G}(\hat{x}) && (25)
\end{aligned}$$

By setting the approximant dynamics to satisfy $\hat{G}(\hat{x}) = G(\hat{x})$, we can set the networks to have the dynamics of our target system. This is achieved with the following optimization problem for ϕ_j:

$$\begin{aligned}
C(\boldsymbol{\phi}) &= \int_X \left(\sum_{j=1}^{N} \boldsymbol{\phi}_j f(\boldsymbol{\eta}_j \cdot \boldsymbol{x} + \beta_j) - \tilde{G}(\boldsymbol{x}) \right)^2 d\boldsymbol{x} + \alpha \boldsymbol{\phi}^T \boldsymbol{\phi} && (26)\\
\Rightarrow \quad & \tilde{G}(\boldsymbol{x}) = \tau_s G(\boldsymbol{x}) + \boldsymbol{x} && (27)
\end{aligned}$$

where $\tilde{G}(\boldsymbol{x})$ is the target function for optimization. One can also resolve the resulting optimization problem for $\boldsymbol{\phi}$ when the spikes are filtered by a double exponential synapse with rise time τ_R and decay time τ_D. The only difference is that the target function changes:

$$\tilde{G}(\boldsymbol{x}) = \boldsymbol{x} + (\tau_R + \tau_D) G(\boldsymbol{x}) + \tau_R \tau_D D_{\boldsymbol{x}} G(\boldsymbol{x}) \cdot G(\boldsymbol{x}) \tag{28}$$

With the weight matrix generated by the NEF approach, one can immediately simulate the network of spiking neurons. While these networks can mimic a dynamical system, they incorporate nothing about the spike timing and neuronal correlations. The accuracy of these networks can be improved upon by including this information.

The key to incorporating spike times and neuronal correlations in the resulting networks is to decouple the weight matrix from the approximant. If we consider our approximant again, then we can rewrite it as the following:

$$\begin{aligned}
\hat{x}(t) &= \boldsymbol{\phi}^T \boldsymbol{r}(t) = \left((\boldsymbol{\eta}^T \boldsymbol{\eta})^{-1} \boldsymbol{\eta}^T \right) \left(\boldsymbol{\eta} \boldsymbol{\phi}^T \boldsymbol{r}(t) \right) && (29)\\
&= \boldsymbol{M} \boldsymbol{\omega} \boldsymbol{r}(t) = \boldsymbol{M} \boldsymbol{s}(t) && (30)
\end{aligned}$$

and so we can write our approximant equivalently as a linear combination of currents, as opposed to a linear combination of firing rates. At first glance, it would seem that the quantity $\boldsymbol{M}$ is also fixed, since altering $\boldsymbol{M}$ would alter the encoders as $\boldsymbol{M} = (\boldsymbol{\eta}^T \boldsymbol{\eta})^{-1} \boldsymbol{\eta}^T$. However, if we consider any other $\tilde{\boldsymbol{M}}$, then this generates a new pair of encoders and decoders $\tilde{\boldsymbol{\eta}}$ and $\tilde{\boldsymbol{\phi}}$ given by:

$$\begin{aligned}
\tilde{\boldsymbol{\eta}} &= (\tilde{\boldsymbol{M}}^T \tilde{\boldsymbol{M}})^{-1} \tilde{\boldsymbol{M}}^T && (31)\\
\tilde{\boldsymbol{\phi}}^T &= \tilde{\boldsymbol{M}} \boldsymbol{\omega} && (32)
\end{aligned}$$

which yields $\tilde{\boldsymbol{\eta}} \tilde{\boldsymbol{\phi}}^T = \boldsymbol{\omega}$. Thus, M serves as a free parameter. Therefore, we can use $\boldsymbol{M}$ to carry out another round of optimization. Now however, we will optimize in time. Consider the following time dependent error:

$$\begin{aligned}
D(\boldsymbol{M}) &= \int_0^T \left(\hat{\boldsymbol{x}}(t') - \boldsymbol{x}(t') \right)^2 dt' + \lambda \boldsymbol{M}^T \boldsymbol{M} && (33)\\
&= \int_0^T \left(\boldsymbol{M} \boldsymbol{s}(t') - \boldsymbol{x}(t') \right)^2 dt' + \lambda \boldsymbol{M}^T \boldsymbol{M} && (34)
\end{aligned}$$

Solving this optimization problem yields a time dependent $\boldsymbol{M}$, which we will relabel $\boldsymbol{M}^*(T)$ and this yields the best approximant the network is capable of generating at time T given the static weight matrix $\boldsymbol{\omega}$. As the weight matrix is static, one can compute $\boldsymbol{M}^*(T)$ for some sufficiently large T such that the correlations generated by the neuronal currents have reached their stationary state. This matrix can be used to generate an approximant for all time $t > T$. The optimization only has to be performed once, at some suitably large time T as opposed to the FORCE method which computes the correlation matrix and updates the decoder

continuously in an online fashion. To our knowledge, this procedure has not been previously applied to NEF generated weight matrices.

With low rank weight matrices, we have found that incorporating spike times has little effect in reducing the error. To create a full rank weight matrix, we will add sparsity to the network. Consider neuron j in the network. The currents going into this neuron should still satisfy the linear encoding requirement:

$$s_i(t) = \boldsymbol{\eta}_i \cdot \boldsymbol{x} = \sum_{j \in \sigma(i)} \omega_{ij} r_j(t) \tag{35}$$

The neuron is now receiving inputs from a subset $\sigma(i)$ of neurons from the rest of the network. We will assume that the degree of sparseness in the network is p, and the sparsity is unstructured. We now have N optimization problems that can be solved in parallel:

$$C_i(\boldsymbol{\omega}_i) = \int_x \left(\sum_{j \in \sigma(i)} \omega_{ij} f(\boldsymbol{\eta}_j \cdot \boldsymbol{x} + \beta_j) - \boldsymbol{\eta}_i \tilde{G}(\boldsymbol{x}) \right)^2 d\boldsymbol{x} + \alpha \boldsymbol{\omega}_i^T \boldsymbol{\omega}_i, \quad i = 1, 2, \ldots N \tag{36}$$

This sets the macroscopic dynamics to approximate the intended dynamics. As the matrix is now full rank, we can simulate the network, carry out the time dependent neuronal optimization and extract a better approximant. A schematic of this entire procedure is shown in Figure S1A in addition to the eigenvalues of the weight matrix in Figure S1D. Note the similarity in eigenvalue distribution to FORCE trained weight matrices.

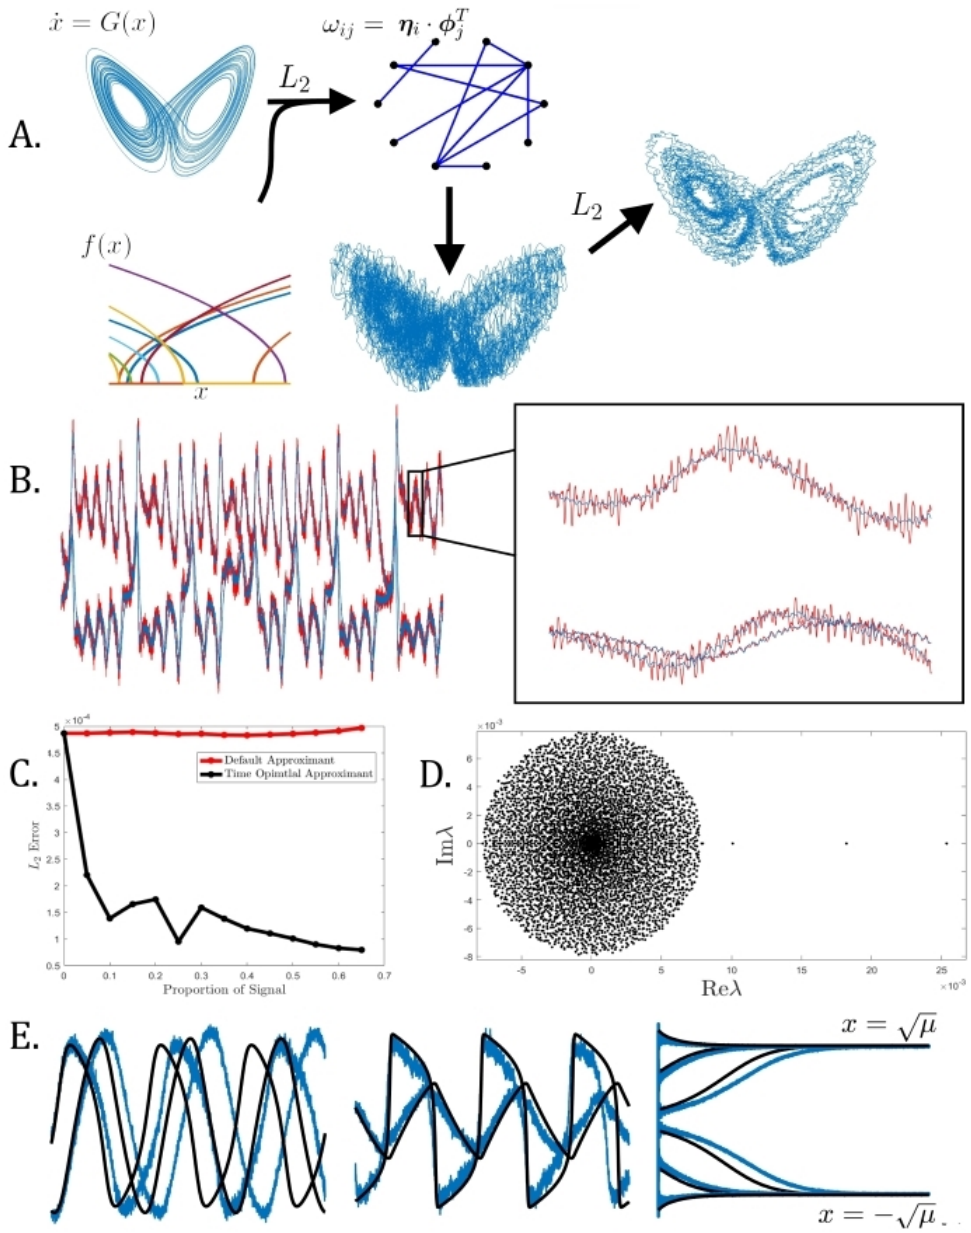

Figure S1: (A) The Neural Engineering Framework (NEF) generates functional spiking neural networks by first specifying a desired set of dynamics and a population of tuning curves over the phase variable x. The assumption that the weight matrix is separable yields a static optimization problem in the L_2 norm such that a linear combination of tuning curves approximates a function $\hat{G}(x)$ that is related to the target dynamics and the specific postsynaptic filter used. Once the spiking neural network is generated, a second L_2 optimization over time can resolve the time optimal approximant after the spiking network is simulated by using either the desired dynamics, or a smoothed version of the network output. (B) The Lorenz system is approximated using the NEF solution for a spiking neural network of 5000 theta neurons coupled together with $p = 0.15$ sparsity. The default NEF approximant is shown in red while the time optimal approximant is shown in blue which uses a smoothed version of the red as a target. (C) As a greater proportion of the signal is used in the approximation, The L_2 error of the remaining proportion (normalized by the remaining signal time) decreases while the L_2 error for the default approximant remains the same. (D) The eigenvalues for the weight matrix used to generate the Lorenz system. When sparsity is added to the network, the eigenvalues become distributed over a circle in the complex plane, with k eigenvalues distributed outside the unit circle for a k dimensional dynamical system. (E) The NEF can generate other dynamics such as the Van der Pol system in the harmonic ($\mu = 1$) or relaxation ($\mu = 5$) phase and the pitchfork topological normal form system with a variety of initial conditions. As the single static L_2 optimization used in the NEF procedure resolves the dynamics over the entire region of phase space, the pitchfork set spiking network exhibits bistability.

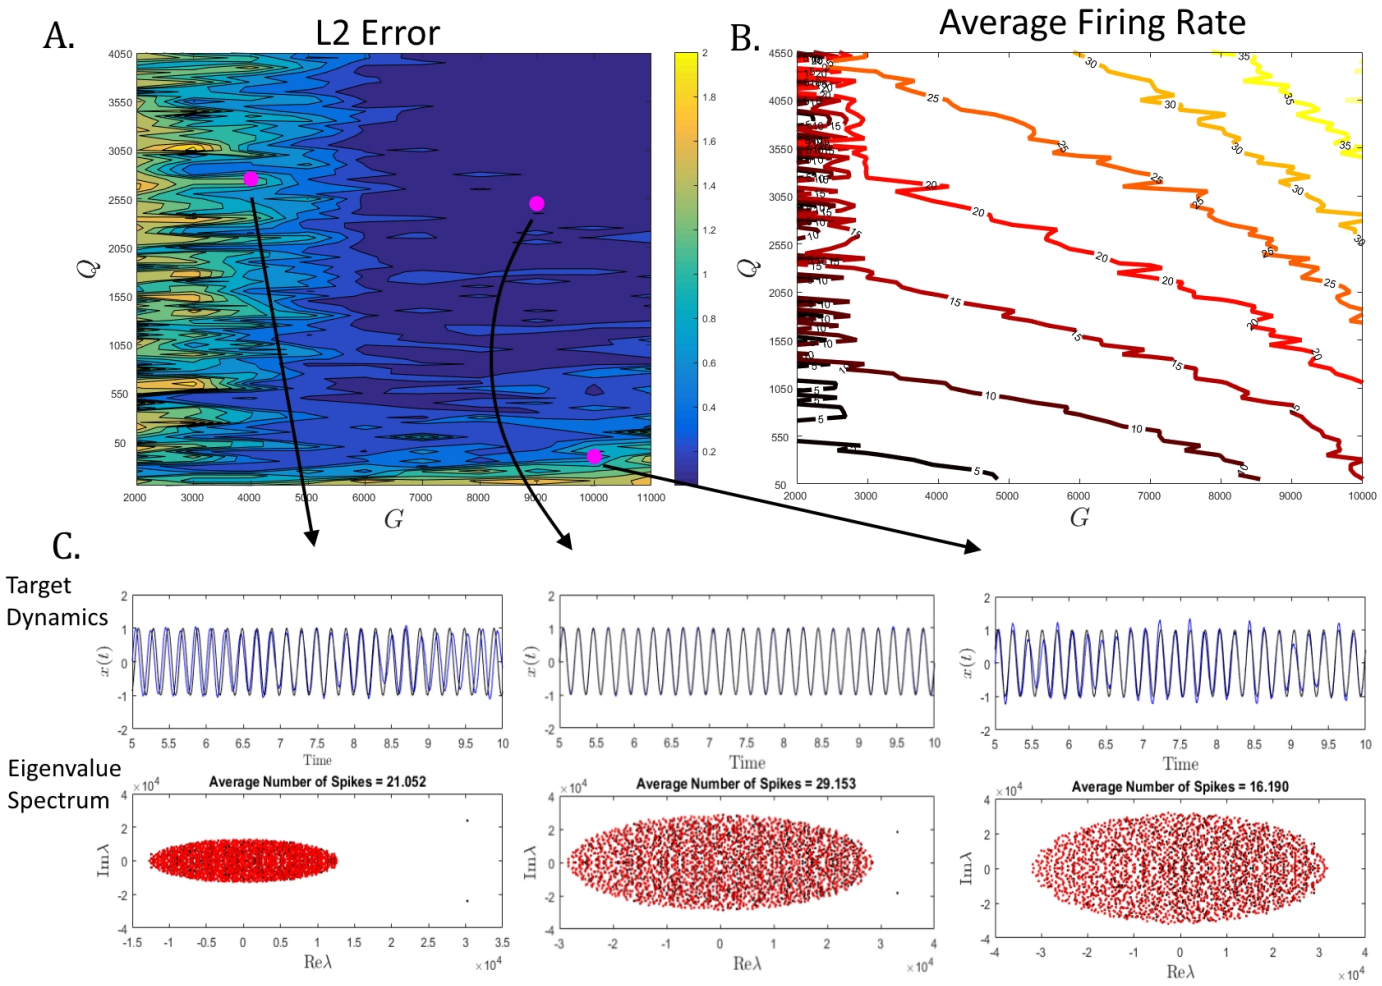

Figure S2: (A) Networks of 2000 Izhikevich neurons were run over a 90×10 uniform mesh over the (Q, G) parameter space. Each network was tasked to learn the dynamics of a 5 Hz sinusoidal oscillation using FORCE training. 4 seconds of FORCE training were used, and the L_2 error was computed for the last 4 seconds where RLMS is turned off. The colors denote the L_2 error with bluer colours indicating less error. (B) The average firing rate for the networks of neurons simulated over the parameter mesh. The majority of these networks have low average firing rates (< 50 Hz). (C) (Top) The target dynamics (black) and the network approximant (blue) for three points in the mesh of simulations corresponding to networks with weight matrices heavily dominant eigenvalues (left), dominant eigenvalues (middle), and no dominant eigenvalues (right). The intermediate region has the lowest L_2 error where there are dominant eigenvalues, but they are not much larger then the cloud in magnitude. Red dots denote the weight matrix before training while black dots correspond to eigenvalues after training.

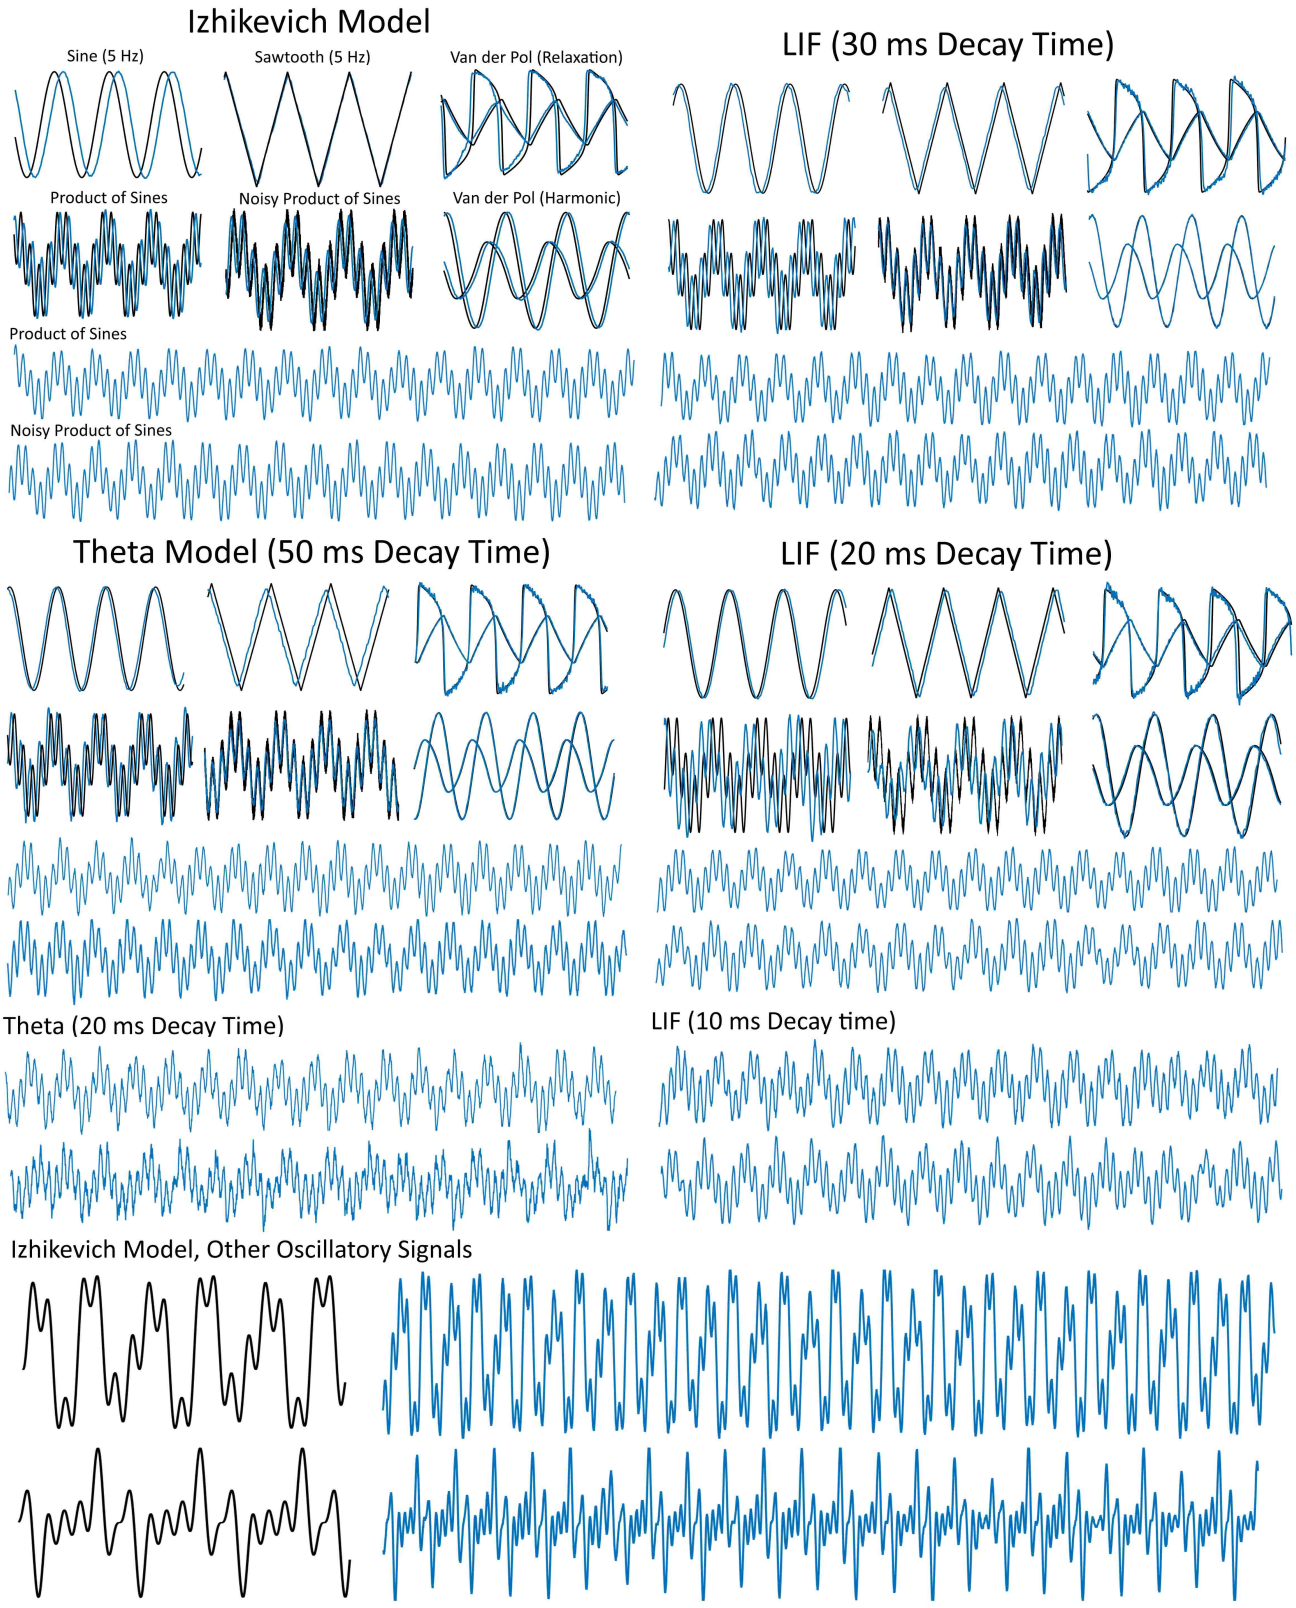
Izhikevich Model
Sine (5 Hz)
Sawtooth (5 Hz)
Van der Pol (Relaxation)
Product of Sines
Noisy Product of Sines
Van der Pol (Harmonic)
Product of Sines
Noisy Product of Sines
LIF (30 ms Decay Time)
Theta Model (50 ms Decay Time)
LIF (20 ms Decay Time)
Theta (20 ms Decay Time)
LIF (10 ms Decay time)
Izhikevich Model, Other Oscillatory Signals

Figure S3: FORCE Trained Oscillators for networks of theta, LIF, and Izhikevich neurons at various time constants. The parameters for each panel can be found in Table S1

Neuron Model	Example	λ^{-1}	G	Q	Training Time	Firing Rate
Izhikevich Model ($\tau_D = 20$ ms)						
	5 Hz sine wave	2 ms	$5 * 10^3$	$5 * 10^3$	5 s	35.7 Hz
	5 Hz sawtooth wave	2 ms	$5 * 10^3$	$4 * 10^3$	5 s	36.8 Hz
	Van der Pol (Harmonic)	2 ms	$1 * 10^4$	$9 * 10^3$	5 s	43.4 Hz
	Van der Pol (Relaxation)	2 ms	$1 * 10^4$	$2 * 10^4$	5 s	41.9 Hz
	$sin(8\pi t)sin(12\pi t)$	2 ms	$1 * 10^4$	$9 * 10^3$	5 s	47.1 Hz
	$sin(8\pi t)sin(12\pi t) + 0.05\zeta_i$	2 ms	$1 * 10^4$	$8 * 10^3$	5 s	47.9 Hz
	$sin(4\pi t)sin(6\pi t)sin(14\pi t)$	2 ms	$1 * 10^4$	$1 * 10^4$	25 s	41.6 Hz
	$\frac{1}{2}sin(8\pi t) + \frac{1}{6}sin(12\pi t) + \frac{1}{4}sin(28\pi t)$	2 ms	$1 * 10^4$	$1 * 10^4$	25 s	52.3 Hz
Theta Model ($\tau_D = 50$ ms)						
	5 Hz sine wave	0.1 ms	50	$2 * 10^4$	5 s	47.53 Hz
	5 Hz sawtooth wave	0.1 ms	50	$2 * 10^4$	5 s	57.65 Hz
	Van der Pol (Harmonic)	0.02 ms	10	10^4	5 s	18.6 Hz
	Van der Pol (Relaxation)	0.02 ms	10	10^4	5 s	13.7 Hz
	$sin(8\pi t)sin(12\pi t)$	0.1 ms	50	$2 * 10^4$	50 s	47.1 Hz
	$sin(8\pi t)sin(12\pi t) + 0.05\zeta_i$	0.1 ms	50	$2 * 10^4$	50 s	47.9 Hz
Theta Model ($\tau_D = 20$ ms)						
	$sin(8\pi t)sin(12\pi t)$	0.01 ms	25	10^4	50 s	47.1 Hz
	$sin(8\pi t)sin(12\pi t) + 0.05\zeta_i$	0.01 ms	15	10^4	50 s	47.9 Hz
LIF Model ($\tau_D = 30$ ms)						
	5 Hz sine wave	0.0025 ms	0.05	10	5 s	47.53 Hz
	5 Hz sawtooth wave	0.0025 ms	0.1	20	5 s	57.65 Hz
	Van der Pol (Harmonic)	0.0025 ms	0.05	30	5 s	18.6 Hz
	Van der Pol (Relaxation)	0.0025 ms	0.05	30	5 s	21.7 Hz
	$sin(8\pi t)sin(12\pi t)$	0.0025 ms	0.1	30	25 s	38.2 Hz
	$sin(8\pi t)sin(12\pi t) + 0.5\zeta_i$	0.0025 ms	0.1	30	25 s	33.9 Hz
LIF Model ($\tau_D = 20$ ms)						
	5 Hz sine wave	0.0025 ms	0.1	30	5 s	38.7 Hz
	5 Hz sawtooth wave	0.0025 ms	0.1	30	5 s	41.9 Hz
	Van der Pol (Harmonic)	0.0025 ms	0.05	30	5 s	25.8 Hz
	Van der Pol (Relaxation)	0.0025 ms	0.05	30	5 s	22.3 Hz
	$sin(8\pi t)sin(12\pi t)$	0.0025 ms	0.1	30	25 s	34.3 Hz
	$sin(8\pi t)sin(12\pi t) + 0.05\zeta_i$	0.0025 ms	0.1	30	25 s	34.7 Hz
LIF Model ($\tau_D = 10$ ms)						
	$sin(8\pi t)sin(12\pi t)$	0.0025 ms	0.1	30	25 s	35.0 Hz
	$sin(8\pi t)sin(12\pi t) + 0.05\zeta_i$	0.0025 ms	0.1	30	25 s	35.0 Hz

Table S1: Parameter table corresponding to figure S3. The quantity ζ_i is an additive white noise signal with mean 0 and standard deviation 1. The integration time step was taken to be 0.01 ms, 0.05 ms and 0.04 ms for the Theta, LIF, and Izhikevich neuron models, respectively.

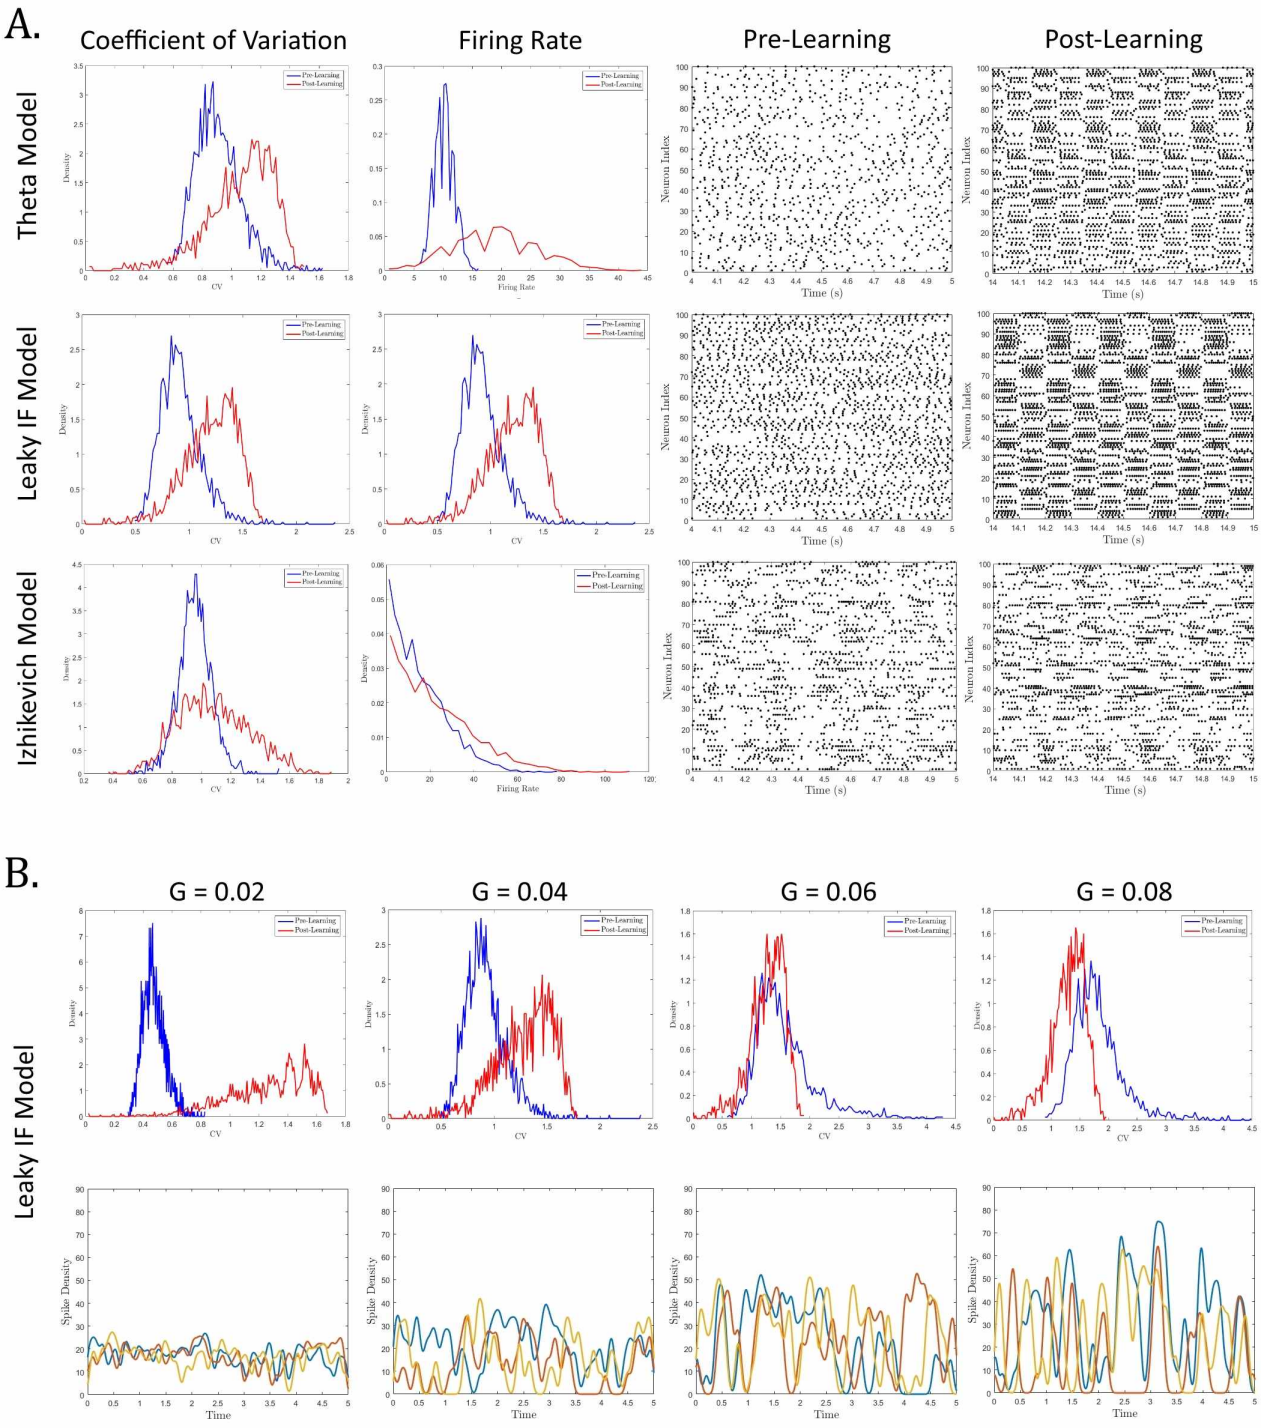

Figure S4: (A) Networks of 2000 theta neurons (top), leaky integrate-and-fire neurons (middle), and Izhikevich models were trained to approximate a 5 Hz sinusoidal oscillator. The coefficient of variation for these networks was centered near 1 prior to learning, indicative of Poisson like firing behavior. After learning, the coefficient of variation is increased due to the regular bursting behavior required to approximate the sinusoidal oscillation. The average firing rates also increase after learning. (B) The LIF network is trained with the FORCE method for increasing values of G. The coefficients of variation increase with increasing G prior to training (top). The coefficients of variation increase for low G after training and decrease for large G. The firing rates of 3 randomly selected neurons are shown for increasing G (bottom). The firing rates are computed by convolving the spike train with 50 ms Gaussian kernels. As G increases, the firing rates start to fluctuate heterogeneously across the network.

S2: FORCE Trained Weight Matrices that Respect Dales Law with Synaptic Boundaries

One potential solution to generate FORCE trained weight matrices that respect Dales law is through synaptic boundaries. Consider a network of N neurons with the first N_E excitatory neurons and $N_I = N - N_E$ inhibitory neurons. We start with static weight matrices ω_{ij}^0 given by the following:

$$\omega_{ij}^0 = \begin{cases} \frac{G}{\sqrt{N}p} & \text{with probability } p \text{ if } i \leq N_E \\ -\frac{G\kappa_i}{\sqrt{N}p} & \text{with probability } p \text{ if } i > N_E \\ 0 & \text{otherwise} \end{cases} \tag{37}$$

where κ_i is a term used to set $\sum_{j=1}^{N} \omega_{ij}^0 = 0$ to initialize chaotic spiking in the reservoir. These static weight matrices immediately respect Dales law and it is sufficient (but not necessary) for the low rank RLMS trained weight matrix to also respect Dales law independently. This can be achieved by setting boundaries on synaptic weights when they do not respect Dales law:

$$\omega_{ij} = \begin{cases} Q\boldsymbol{\eta}_i\boldsymbol{\phi}_j & \text{if} \quad 0 \leq \boldsymbol{\eta}_i\boldsymbol{\phi}_j, \quad i \leq N_E \\ Q\boldsymbol{\eta}_i\boldsymbol{\phi}_j & \text{if} \quad 0 \geq \boldsymbol{\eta}_i\boldsymbol{\phi}_j, \quad i > N_E \\ 0 & \text{otherwise} \end{cases} \tag{38}$$

This is done dynamically as the decoders $\boldsymbol{\phi}_j$ are determined through RLMS. In this way, the decoders and the approximant are determined in precisely the same way as before. The only difference is that the approximant interacts with the reservoir in a manner which obeys Dales law. The results of this procedure are shown in Figure S5 for a network of 2000 Izhikevich neurons with identical parameters learning a 5 Hz sinusoidal oscillator.

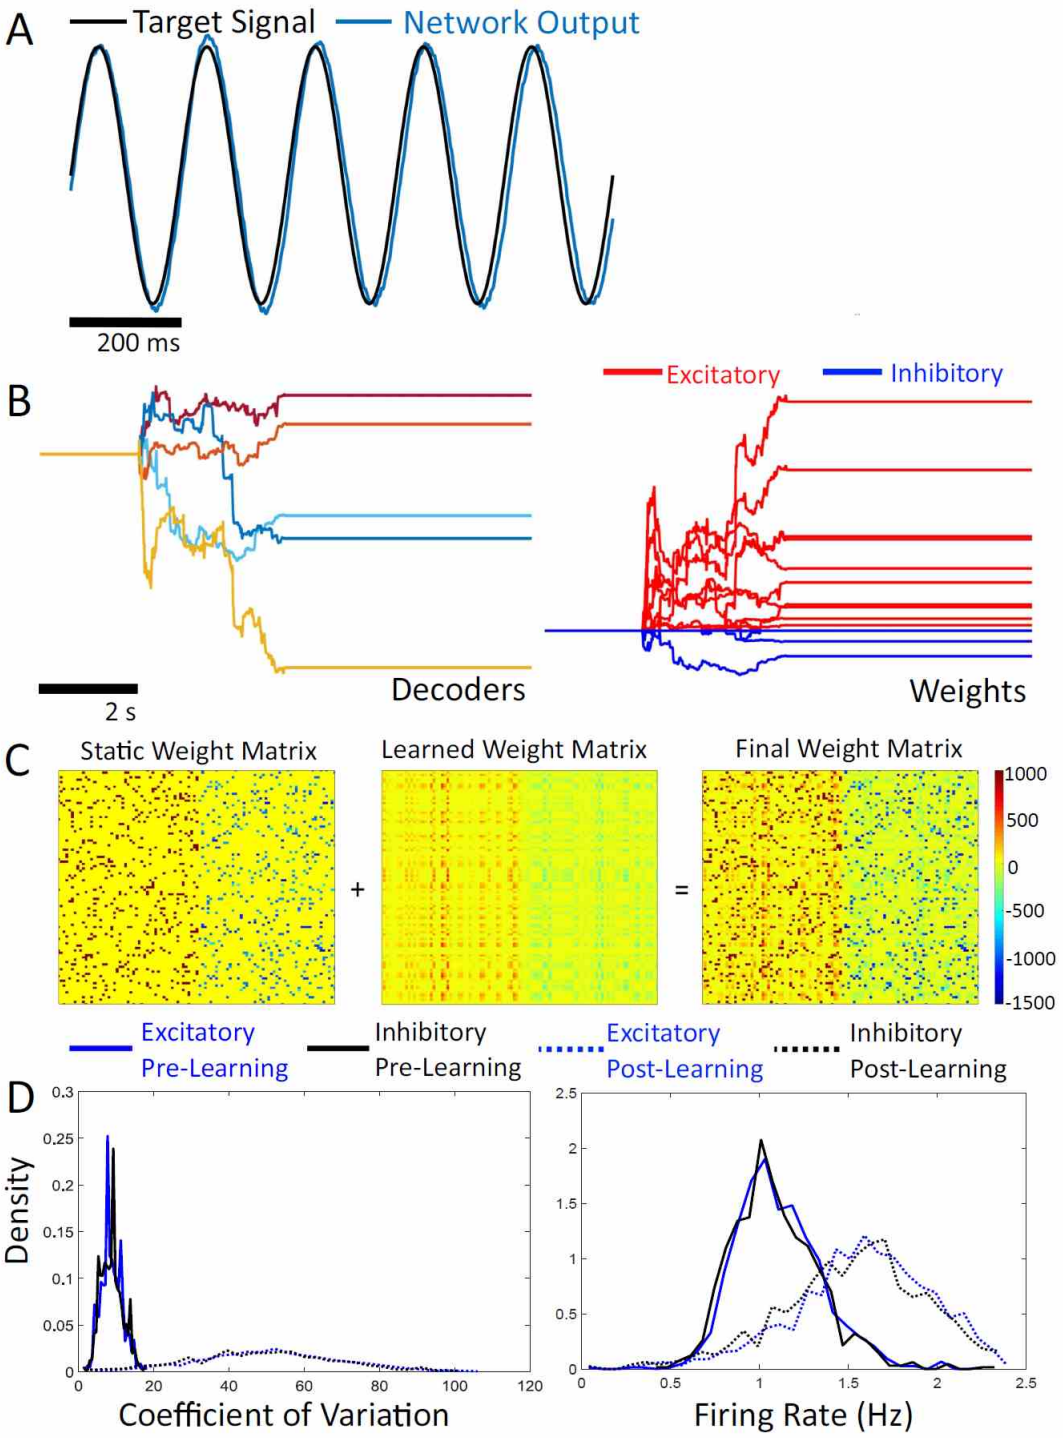

Figure S5: (A) A network of 2000 Izhikevich neurons is trained to mimic a 5 Hz sinusoidal oscillation. The first 1000 neurons are excitatory while the next 1000 are inhibitory. The weight matrices are determined such that both the static weight matrix and the learned weight matrices independently respect Dale's law, and thus the final weight matrix also respects Dale's law. (B) The decoders for the network are learned with 3 seconds of FORCE learning after a 2 second transient, and are not constrained by signs. (C) The time varying component of the weight matrix is given by $\omega_{ij} = \boldsymbol{\eta}_i \boldsymbol{\phi}_j$ when the weight respects Dale's law, and 0 otherwise. (D) The static weight matrix forms a backbone of strong connections that respect Dales law, while the learned weight matrix forms a set of weaker connection weights that force the network to perform the desired dynamics. The final weight matrix is the sum of these two and also respects Dale's law. The final weight matrix had a 0.45 degree of sparseness. (E) The spiking statistics for both excitatory and inhibitory neurons are largely identical to prior work with weight matrices that did not respect Dales law. The spiking prior to FORCE learning is near Poissonian with a coefficient of variation distribution that centers a 1 and increases post-learning. The firing rates also increase post learning with an average of 50 Hz for both populations.

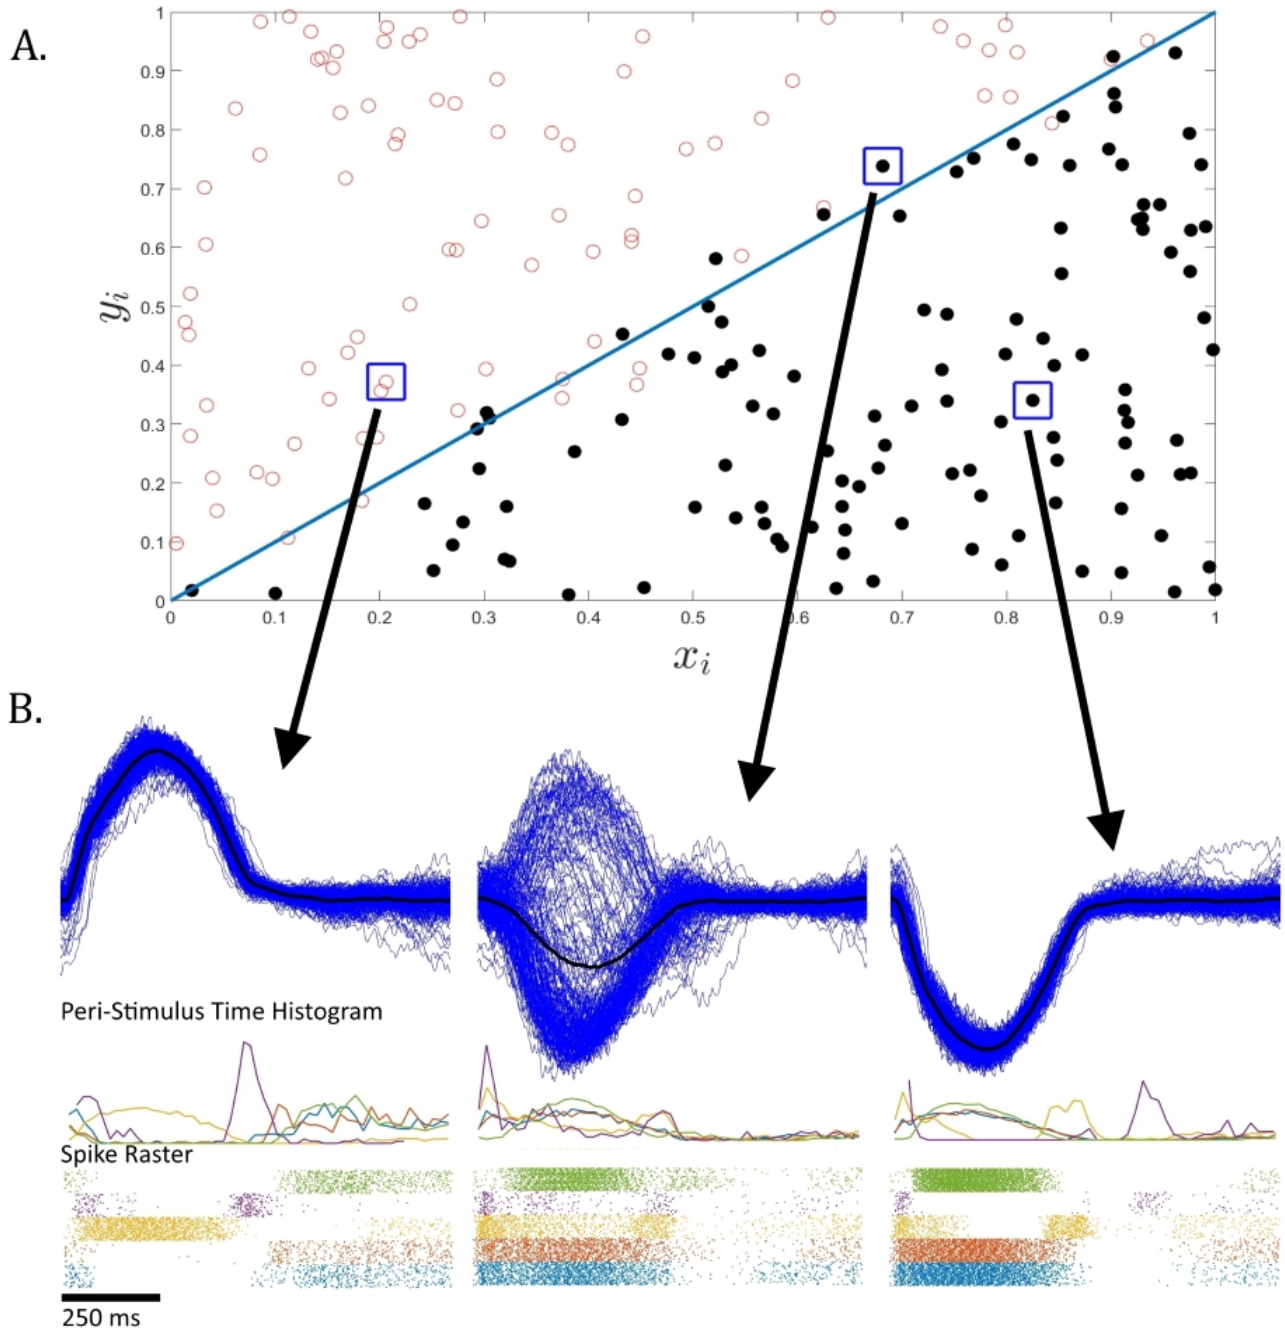

Figure S6: (A) The inputs are classified into two classes by the FORCE trained classifier. Up-states correspond to one class (open red circles) while down states correspond to another class (closed black circles). The boundary between classes is given by $y = x$. (B) The three inputs corresponding to the three arrows from figure (A) are repeatedly displayed for the network to classify 50 times. A peri-stimulus time histogram is generated for the same 5-neurons in each three cases using the raster plots. The network misclassifies points near the boundary (blue line) by displaying an incorrect up/down-state. The misclassification depends on the initial condition of the network and points closer to the boundary are more likely to be misclassified.

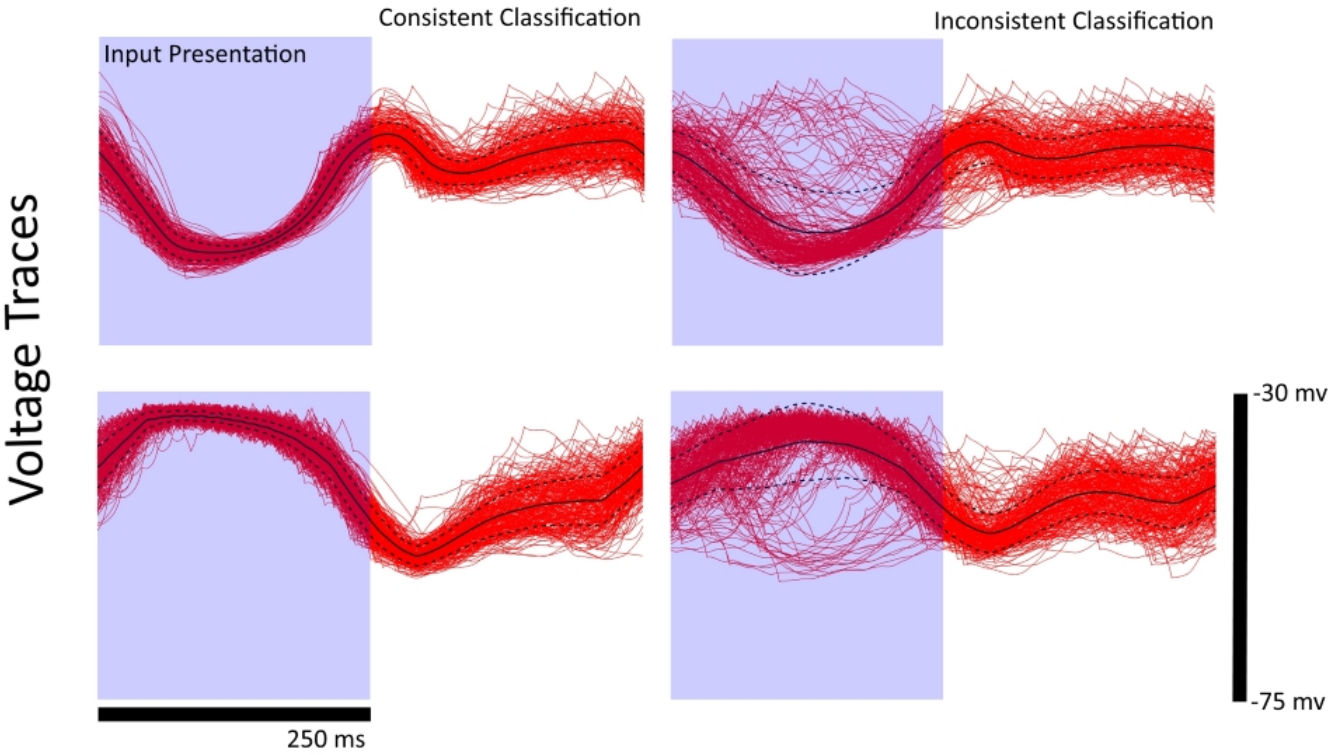

Figure S7: The network of neurons is presented with an identical input 250 times periodically that is either consistently classified (far from the classification boundary) or is inconsistently classified (near the classification boundary). The voltage for a pair of neurons is box filtered with a filter of 80 ms to remove spikes. During input presentation, the voltage variance decreases regardless of whether or not the mean voltage increases (bottom) or decreases (top) for consistent classifications while the voltage variance increases with input presentation when the network inconsistently classifies a point.

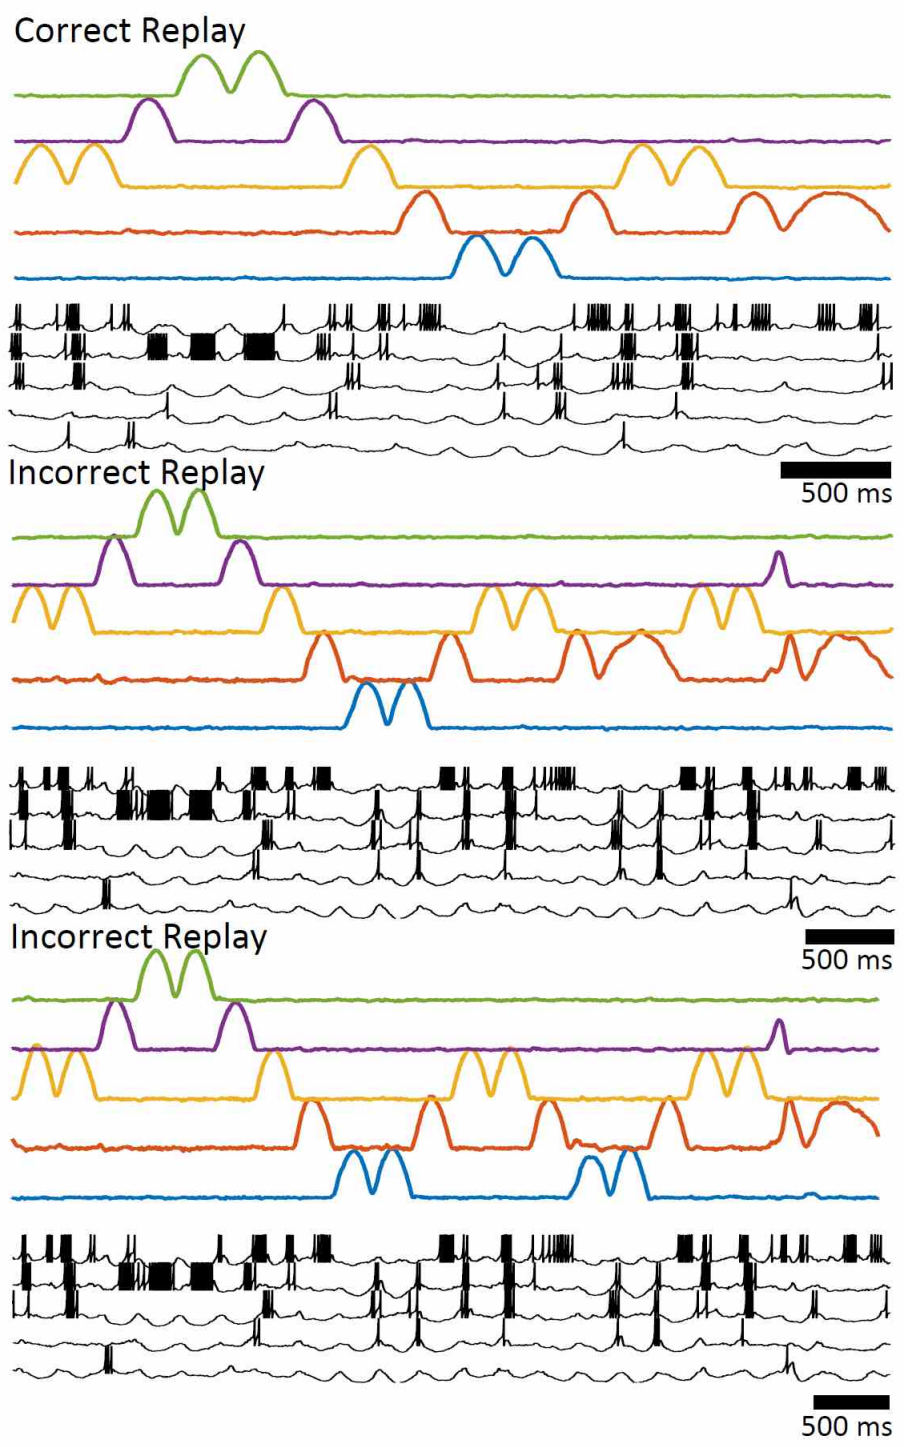

Figure S8 A pair of incorrect repetitions of Ode to Joy after RLMS is turned off (middle, bottom) in addition to a correct repetition (top). The corresponding voltage traces for 5 randomly selected neurons are shown. In both cases, the network fails to elicit the correct note in the repetition, and instead replays an alternate part of the song. The note prior to the mistake heavily influences which part of the song is replayed.

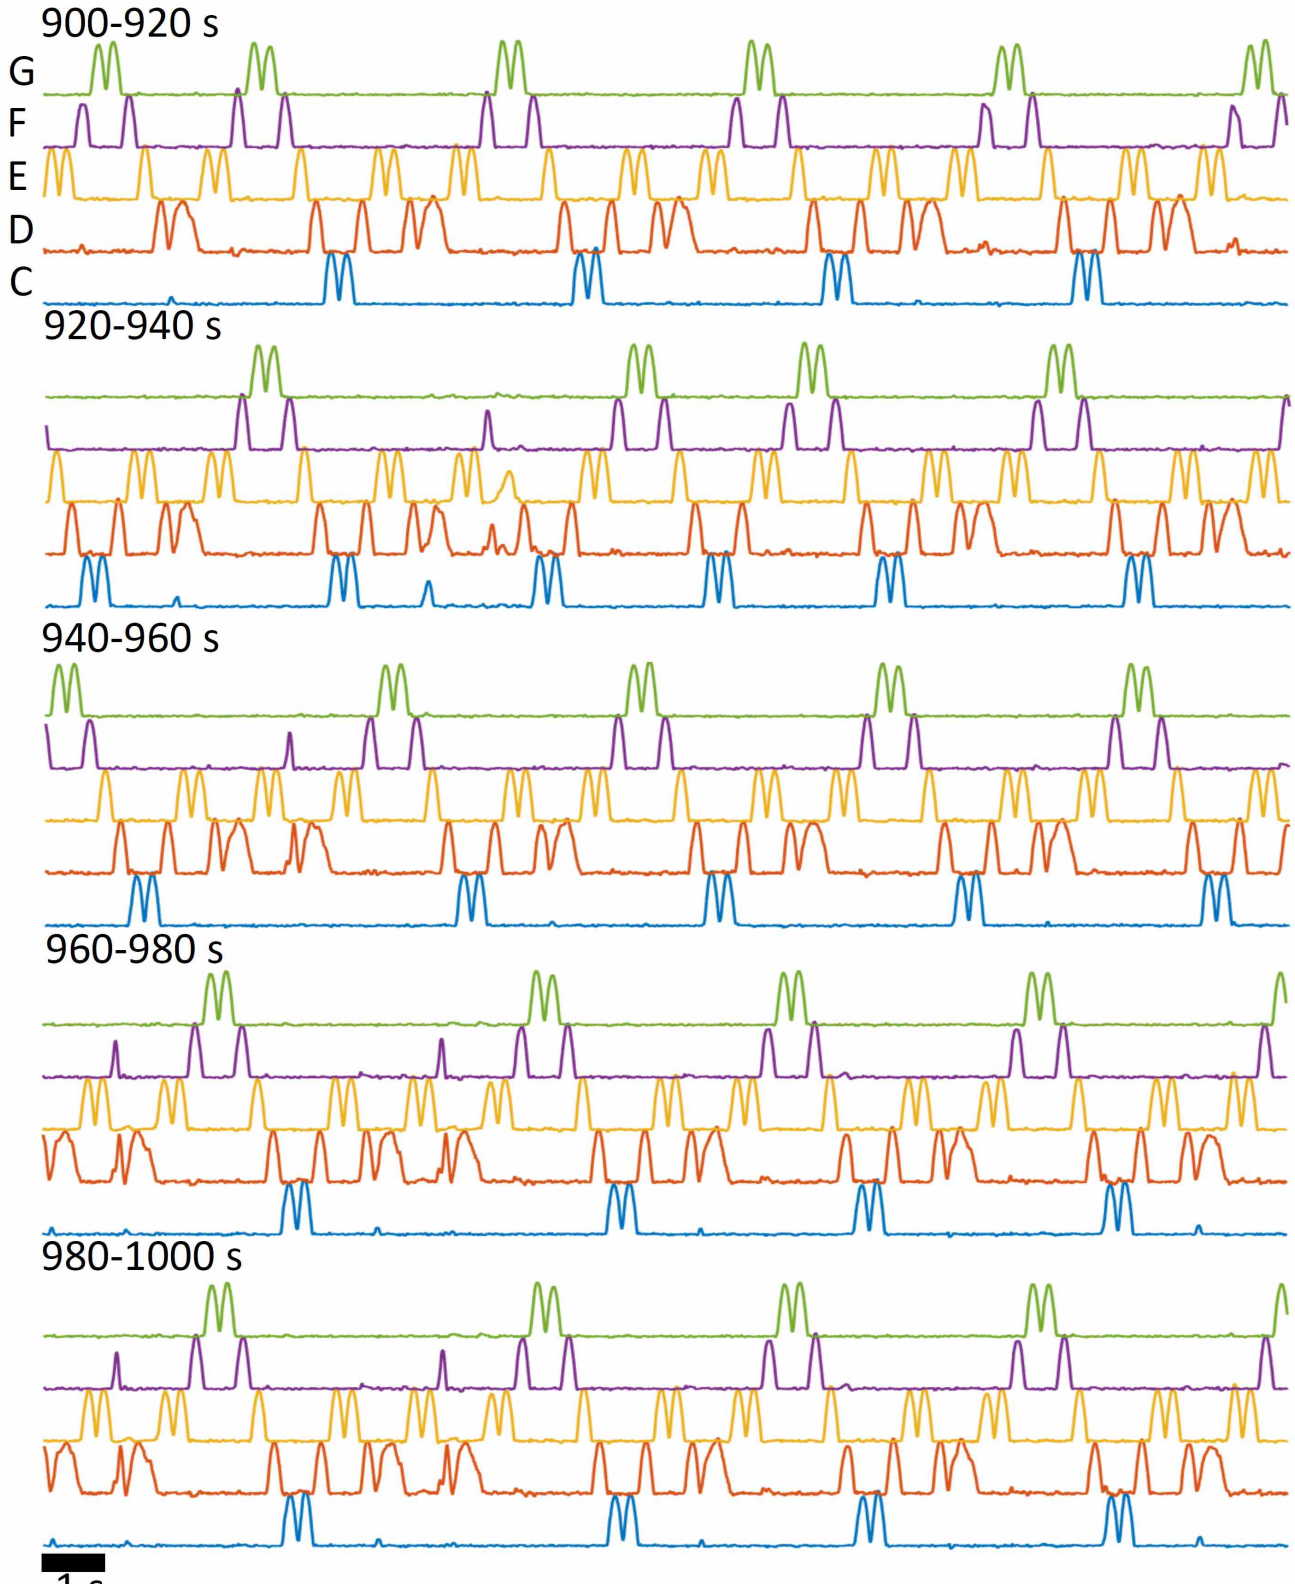

Figure S9: Shown above is 100 seconds of the network output after RLMS learning is turned off for the song example.

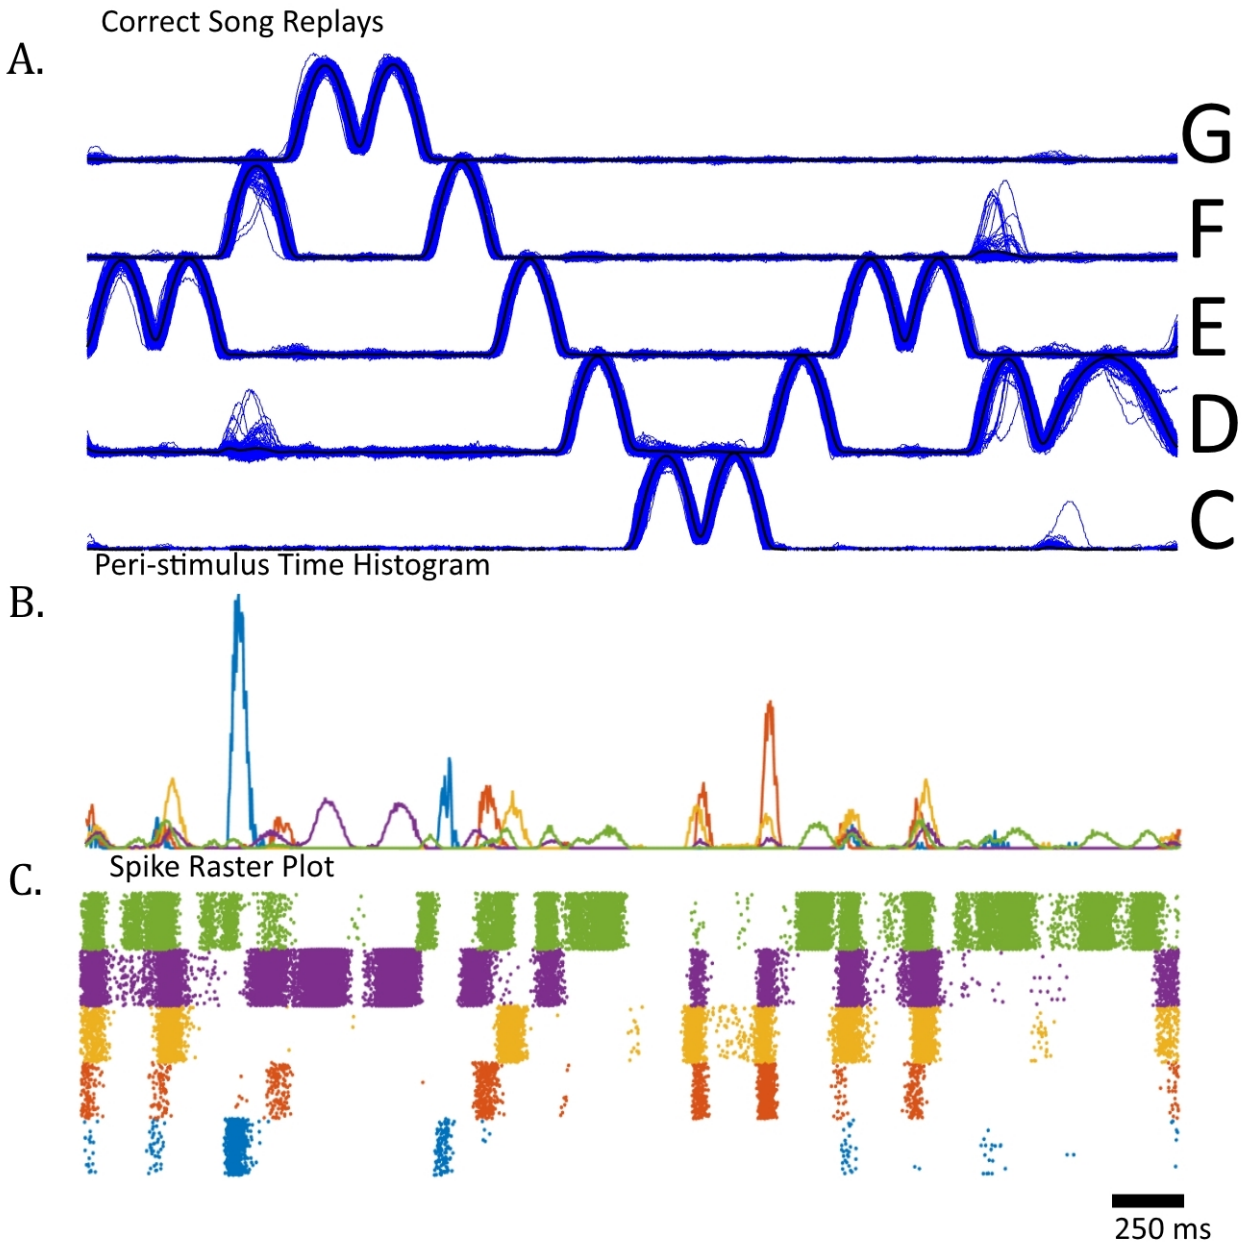

Figure S10: (A) The network continuously replays the song for 1000 seconds after FORCE method is turned off. Correct replays are automatically classified using the L_2 error with the teaching signal and 205 correct repetitions (shown in blue) are found in the 1000 second stretch corresponding to 82% of the signal. The mean of these repetitions is shown in black. (B) The peri-stimulus time histogram is constructed from the raster plots in (C) of 5 neurons. As shown in the raster plot, there is considerable trial-to-trial variability.

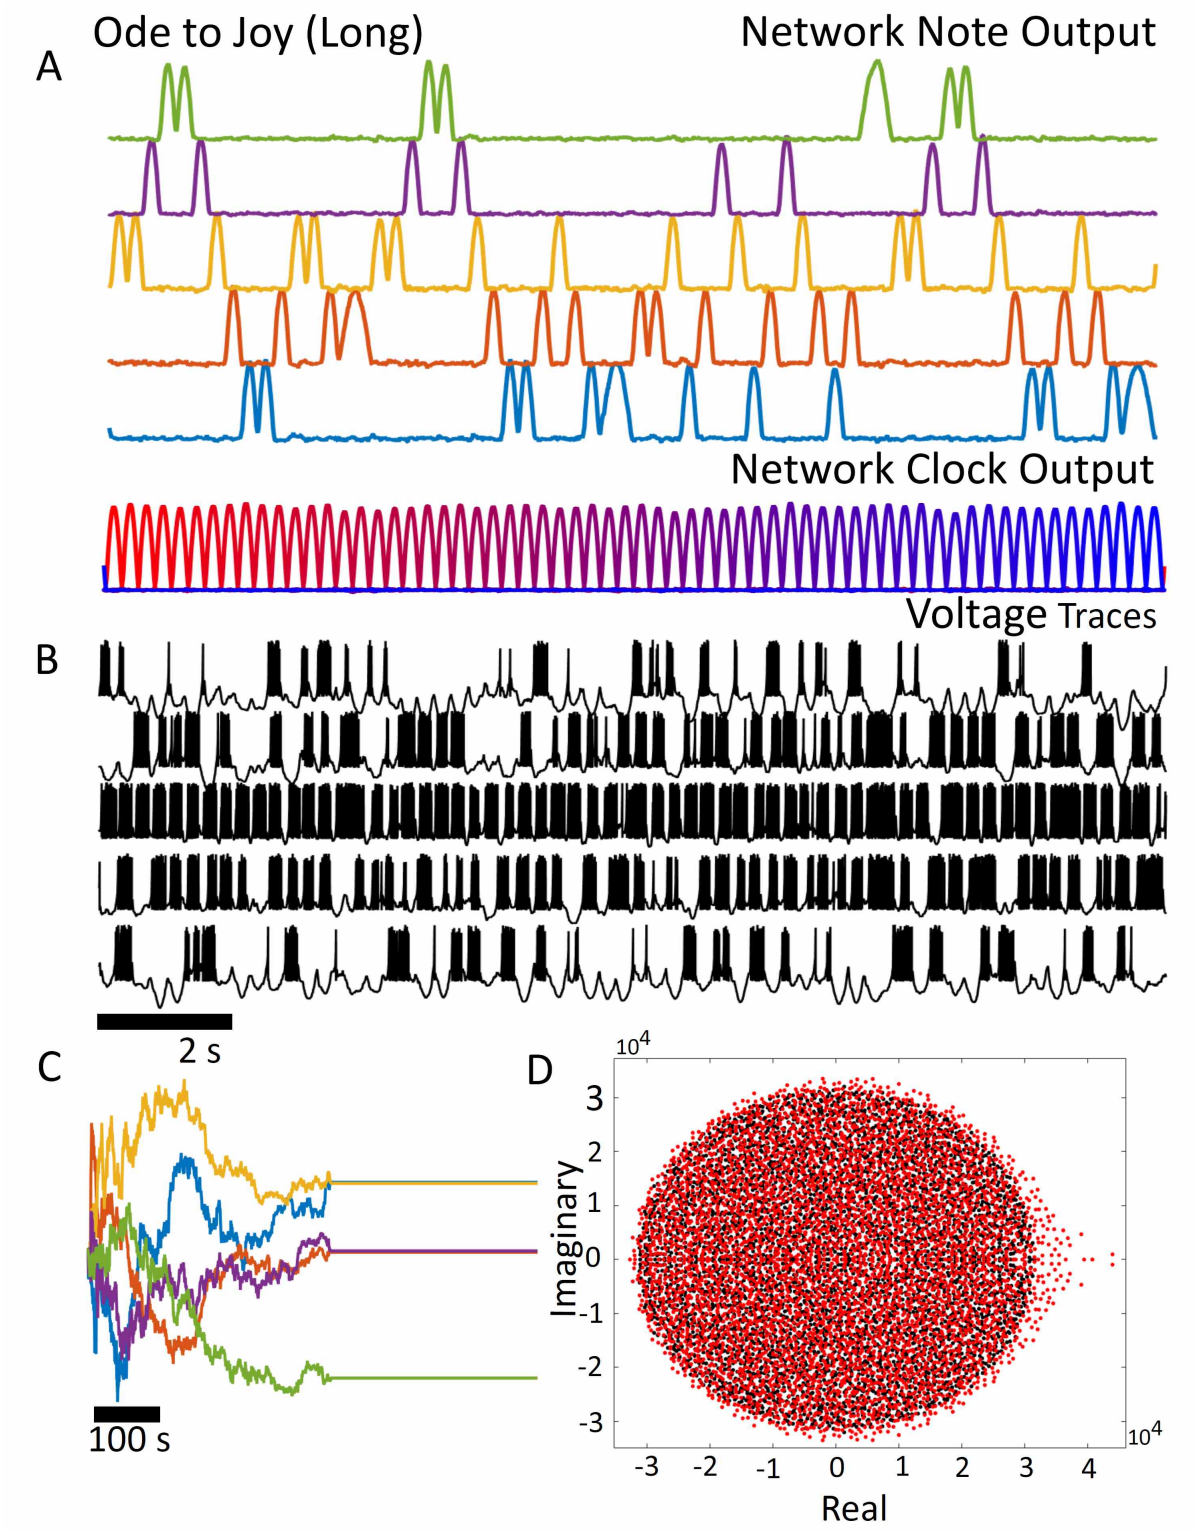

Figure S11: (A) A network of 5000 Izhikevich neurons correctly learns the 64 note sequence to the song Ode to Joy in addition to a 64 component sequence of upstates that correspond to the internal network clock. Unlike other implementations, the network clock sequence corresponds to the 6th-69th components of the supervisor presented to the network with the red blue colour spectrum indicating the position of the upstate in the sequence (B) The voltage traces for 5 neurons are shown. Unlike the previous implementation we considered, the neurons in this network encode information about time through the final 64 components and the network note through the first 5 components of the encoders and decoders. (C) The network was trained with FORCE training for a period of 379 seconds, after which RLMS was deactivated. The network performed the correct 64 note sequence for the remaining 321 seconds of the simulation. Unlike our previous implementation, recall of the entire song was flawless for the remainder of the simulation. (D) The eigenvalue spectrum of the weight matrix after learning shows a cloud of eigenvalues (red) near the static weight matrices original circle of eigenvalues (black) after learning this high dimensional supervisor.

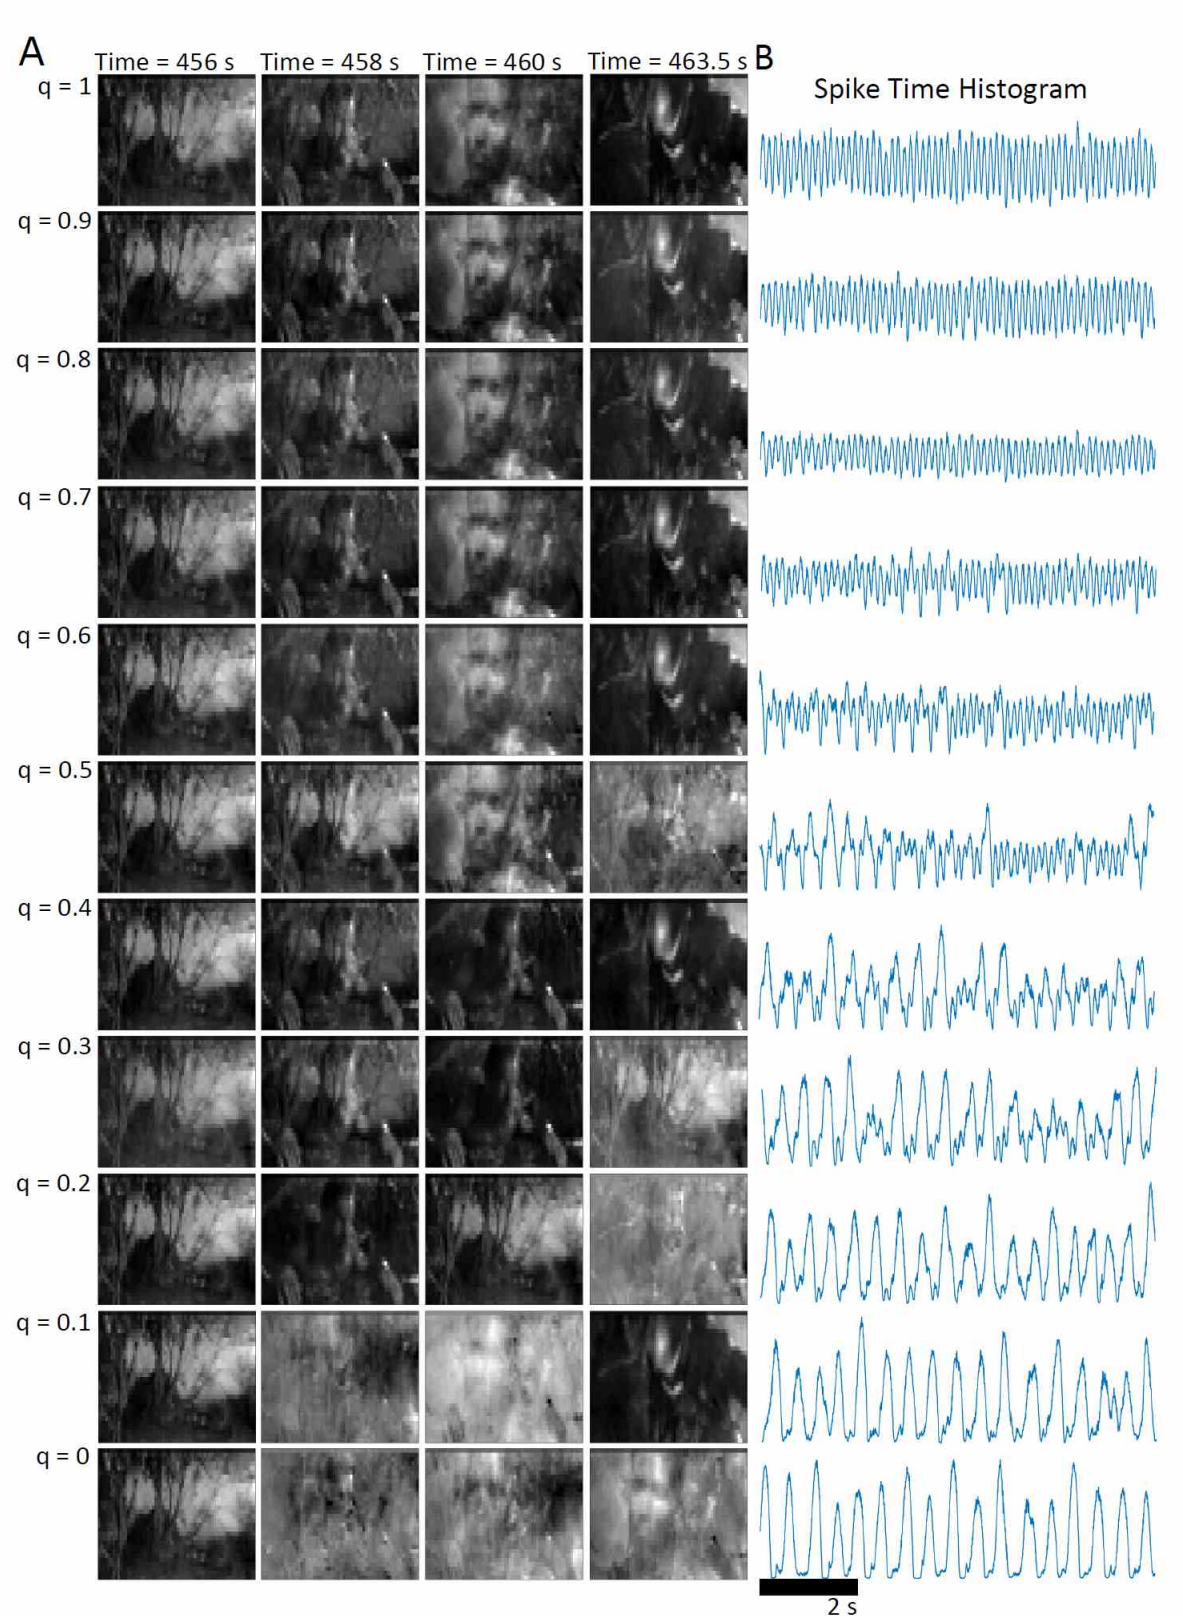

Figure S12: (A) The movie is played by the network and displayed at different time points after training for decreasing levels of clock amplitude. As the amplitude of the clock input is decreased, the recall performance decreases. (B) As the clock amplitude decreases, the power of the theta oscillation in the local field potential decreases, eventually transitioning to a delta oscillation.

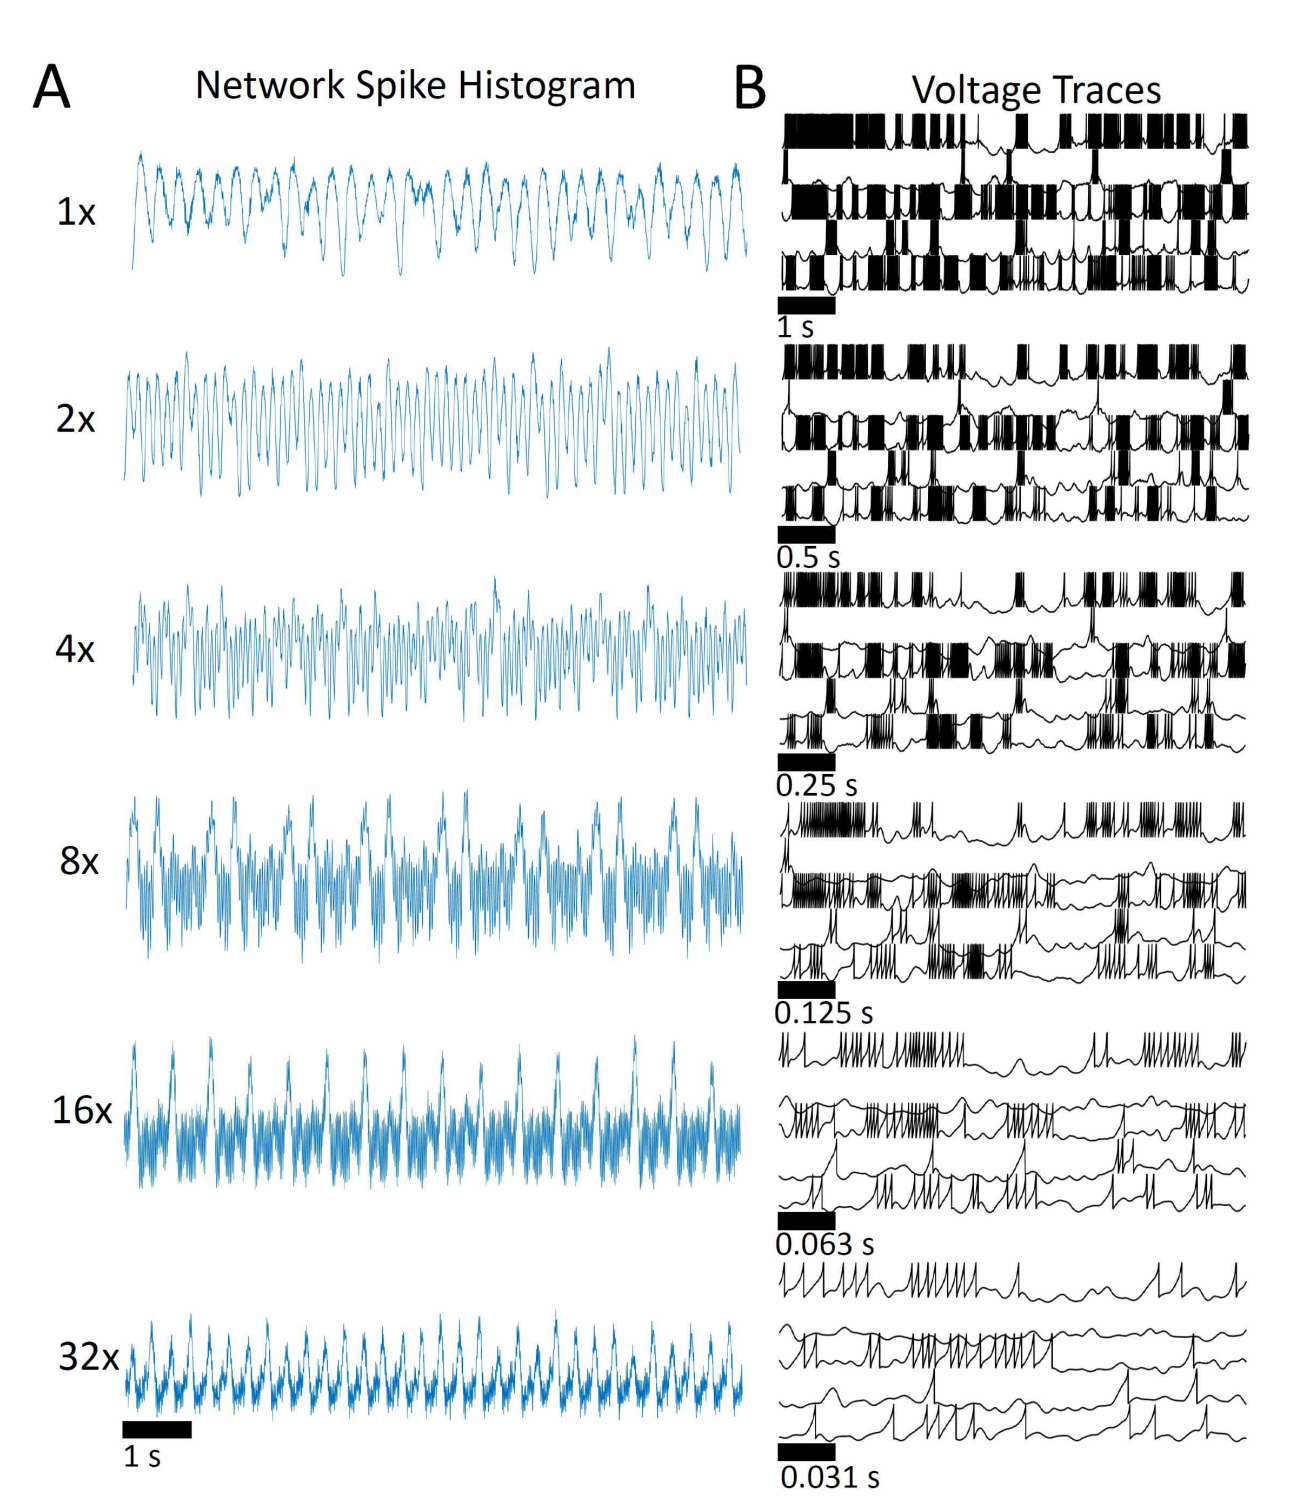

Figure S13: (A) The spike-time histogram for the entire network is computed for different levels of clock compression. As the compression factor increases, the histogram has higher frequency components. The histogram is computed with 5 ms time bins. (B) The voltage traces for 5 randomly selected neurons for a single replay of the movie scene at different compression ratios. The traces are compressed by a similar ratio as the LFP, and the movie replay, however the relative firing rate decreases with increasing compression.